

Deformable Models for Volume Feature Tracking

by
Gregory James Klein

B.S. (University of Nebraska, Lincoln) 1985
M.S. (University of Maryland, College Park) 1989

A dissertation submitted in partial satisfaction of the requirements for the degree of
Doctor of Philosophy
in
Vision Science
in the
GRADUATE DIVISION
of the
UNIVERSITY of CALIFORNIA at BERKELEY

Committee in charge:

Professor Jitendra Malik, Chair
Professor Stanley Klein
Associate Professor David Forsyth

Spring 1999

The dissertation of Gregory James Klein is approved:

Chair

Date

Date

Date

University of California, Berkeley

Spring 1999

Deformable Models for Volume Feature Tracking

Copyright © 1999
by
Gregory James Klein
All rights reserved

Abstract

Deformable Models for Volume Feature Tracking

by

Gregory James Klein

Doctor of Philosophy in Vision Science

University of California, Berkeley

Professor Jitendra Malik, Chair

Conventional three dimensional medical imaging devices have made possible the routine visualization of biological tissue acquired as a sequence of image volumes over time. Because biological tissue is rarely static, the accurate registration of features in two volumes from an image sequence is a common requirement for analysis of these data. This registration task can be difficult, since tissue can bend and stretch over time, and a motion description capturing the non-rigid deformation can be quite complex. One way to characterize the non-rigid deformation is by means of a vector field called a motion field, which describes the relative displacement of each voxel, and thus establishes a correspondence between any set of features in the two volumes. A field such as this can adequately describe any non-rigid deformation seen in biological tissue; however, algorithms designed to estimate it are often confounded by the large dimensionality of the problem. Given the size of conventional medical imaging data sets, the total degrees of freedom represented by the number of independent vectors in the motion field is tremendously large.

This dissertation focuses on incorporating elastic material models into a motion estimation algorithm so that the deformation of tissue in a medical imaging data set can be more accurately described. The research has been motivated by a problem seen in the acquisition of cardiac Positron Emission Tomography (PET). Because the heart moves during a PET acquisition, data are often distributed into different time frames, each capturing a specific phase of the cardiac cycle. This reduces motion-induced blur; however, the individual time frames are quite noisy, and need to be recombined in some manner to improve image quality. Since the shape of the heart is different in each time frame, the data

may not be simply added together, but first must be warped to match a common reference shape. Warping is achieved by using an optical flow-based motion algorithm to estimate the mapping between corresponding voxels in a source and a reference volume. The material model serves to reduce the large domain of possible motion fields, constraining the image data to deform as if it were a physical piece of elastic media undergoing stress from an external force.

The main result of this dissertation is that by better modeling the material properties of tissue within the field of view in a dynamic cardiac PET acquisition, a better estimation of the motion field describing the deformation can be obtained. The motion estimation algorithm is unique from past approaches in that it uses a non-uniform model allowing large-displacement deformations to describe the elastic properties of a cardiac volume. It also uses a forward sampling scheme appropriate for recombination of voxels into a composite motion-corrected volume. Results indicate that this motion field can be used to produce a composite data set with less motion blur and improved contrast to noise characteristics.

In memory of Dad

Table of Contents

CHAPTER 1. Introduction.....	1
1.1. Introduction	1
1.2. Contribution of this Dissertation	4
1.3. Summary of Results	6
1.4. Dissertation Outline.....	7
CHAPTER 2. Related Work	9
2.1. Deformable Model Applications	9
2.2. Optical Flow Based Work	11
2.3. Simple Regularization Constraints	13
2.4. Miscellaneous Cardiac Motion Tracking Techniques	13
2.5. Perspectives on Related Work.....	14
CHAPTER 3. Imaging Model	16
3.1. Tomography - PET Image Reconstruction.....	17
3.2. Gated PET Acquisitions	18
3.3. Imaging Model Implications	21
CHAPTER 4. General Formulation	22
4.1. Deformation Framework	22
4.2. Similarity Measures.....	24
4.3. Regularization Constraints	27
4.4. Criteria Minimization	31
4.4.1. Euler-Lagrange Approach.....	31
4.4.2. Nonlinear Over-Relaxation Method	33
CHAPTER 5. Deformation Sampling	35
5.1. Forward Verses Backward Sampling	35
5.2. Sampling Examples	38
CHAPTER 6. Anisotropic Elastic Model	44
6.1. Stress and Strain Relations	44
6.2. Strain Energy Functions	50
6.3. Cardiac PET Piece-wise Isotropic Strain Energy Model	53
CHAPTER 7. Final Formulation	56
7.1. Algorithm Summary.....	56
7.2. Multi-Resolution Approach.....	59
7.3. Calculating the Composite Volume	60

CHAPTER 8. Results	61
8.1. Simulation Models	61
8.1.1. MCAT Phantom	61
8.1.2. Prolate Spheroid Finite Element Model	63
8.2. Material Model Parameter Selection for Isotropic Materials	64
8.2.1. Convergence Properties	69
8.3. Small Displacement Verses Large Displacement Model Effects	70
8.4. Sampling Effects	73
8.5. Piece-wise Isotropic Strain Energy Results	74
8.6. Motion Compensation Improvement	79
8.7. Results using Real Data	81
CHAPTER 9. Conclusions	86
9.1. Discussion	86
9.2. Future Work	89
APPENDIX A. Minimization Via the Calculus of Variations	92
A.1. Euler-Lagrange Equations	92
A.2. Conjugate Gradient Solution	94
Bibliography	97

List of Figures

CHAPTER 1.	1
Figure 1-1. Cardiac Motion Compensation Scheme	3
CHAPTER 2.	9
CHAPTER 3.	16
Figure 3-1. Typical FDG Cardiac PET Image	17
Figure 3-2. Ungated Cardiac Data	19
Figure 3-3. Gated Cardiac Data	20
CHAPTER 4.	22
Figure 4-1. Voxel Correspondence	23
CHAPTER 5.	35
Figure 5-1. Deformation Sampling Techniques	36
Figure 5-2. Ellipsoidal Phantom	38
Figure 5-3. Simple Translation	39
Figure 5-4. Uniform Rotation	40
Figure 5-5. Incompressible Stretch Deformation	42
CHAPTER 6.	44
Figure 6-1. Stress Tensor	44
Figure 6-2. Volume Element Undergoing Deformation	46
Figure 6-3. Simple Shear Example	48
CHAPTER 7.	56
CHAPTER 8.	61
Figure 8-1. MCAT Cardiac Phantom	62
Figure 8-2. Parametric Finite Element Model	63
Figure 8-3. Hollow Ellipsoid Model	64
Figure 8-4. Ellipsoid Incompressible Deformation	65
Figure 8-5. Ellipsoid Compressible Deformation	66
Figure 8-6. Strain Energy for Incompressible Stretch	69
Figure 8-7. Algorithm Convergence Properties	70
Figure 8-8. Strain Energy Model Comparison	71
Figure 8-9. Strain Energy Model Comparison - FEM Phantom	72
Figure 8-10. Sampling Effects	73
Figure 8-11. Motion Flow Error Due to Material Mismatch	75

Figure 8-12. Isotropic vs. Piece-wise Isotropic Model - Noise Free.....	76
Figure 8-13. Isotropic vs. Piece-wise Isotropic - Noisy Case	77
Figure 8-14. MCAT - Isotropic vs. Piece-wise Isotropic	78
Figure 8-15. MCAT - Motion Compensation Example.....	80
Figure 8-16. Human PET Data - Warping End Systole	81
Figure 8-17. Human PET Data - Summing Comparison	82
Figure 8-18. Human PET Data - Second Subject.....	83
Figure 8-19. Human PET Data - Summing With Less Smoothing	85
CHAPTER 9.	86
APPENDIX A.	92

Acknowledgments

I am deeply indebted to the staff of the Center for Functional Imaging at the Lawrence Berkeley National Laboratory for their help, especially to Ron Huesman, who provided many insightful comments and whose meticulous attention to detail helped me to better understand issues that I would have otherwise overlooked. I would also like to thank my advisor, Jitendra Malik, whose guidance throughout my program at Berkeley has been extremely helpful, and whose expertise in computer vision techniques has been an inspiration. Special thanks to my good friends for encouraging me to be less sick of my research when I'd rather be racing my bike. Finally, thanks to Pat, my wife, for her patience and support.

Research for this thesis was supported in part by the National Heart, Lung and Blood Institute of the U.S. Department of Health and Human Services under grant HL25840 and in part by the Director, Office of Energy Research, Office of Health and Environmental Research, Medical Applications and Biophysical Research Division of the U.S. Department of Energy under Contract no. DE-AC03-76SF00098.

Chapter 1

Introduction

1.1. Introduction

Positron Emission Tomography (PET) is a three-dimensional (3D) imaging technique which allows *in vivo* visualization of a biologically active radiolabeled tracer distribution within the body. Analysis of cardiac function is a frequent application of PET. Determining the viability of cardiac tissue in a compromised heart via metabolic measures or blood flow, measuring heart wall motion or cardiac contractile efficiency, and characterization of the cardiac sympathetic nervous system are all possible means of studying the heart and characterizing its health with PET. However, just as an image may be blurred by an optical camera photographing a moving scene with too slow a shutter speed, so too may cardiac images acquired via PET be blurred by motion due to the fact that the typical time to acquire a cardiac PET image is much longer than the time for one cardiac cycle. This motion blur degrades the diagnostic potential of cardiac PET data.

Peak spatial resolution of conventional whole body PET scanners is nearly 3 mm full width at half maximum (FWHM); however, motion of the heart due to the cardiac cycle can be considerably larger than that distance. Between end diastole and end systole, the base of the heart typically moves 9-14 mm towards the apex, and the walls of the left ventricle thicken at end systole by about 6 mm from an end diastolic thickness ranging from 7.5 to 9.3 mm in a normal human [77,67,53]. Blur due to cardiac motion is therefore one of the principal factors limiting high spatial resolution of cardiac PET image features.

One method to combat this motion blurring is to use a type of data acquisition called *gated* cardiac PET. Here, the cardiac cycle is monitored using an electrocardiograph (ECG) during the acquisition, and tomograph data are directed into different storage locations based on the time since the last R-wave occurrence. By summing these events into the same time frames with respect to the R-wave over all heart beats in an acquisition, a group of reconstructed time frames, or gates, can be obtained which capture the motion of the heart as a set of “freeze frames.” An unfortunate effect of distributing the data into many different time frames is that the statistical quality of the reconstructed volume suffers. Though the gating has removed most of the motion blur, the individual gates are so noisy that their diagnostic quality is often no better than the blurry ungated data. One way to improve the statistical quality of the reconstructed images in a gated study is simply to increase the total acquisition time. In this way, the number of tomographic events in each gate can be made equal to the total number of events in a nongated study. Unfortunately, this tactic is usually infeasible because of short-lived imaging radiotracers or because of the patient’s inability to remain stationary in the scanner for extended periods of time. An alternate approach for improving image statistics is to sum the individual gates to form a composite image, but in order to do so without reintroducing the motion-induced blur, a correspondence must first be found between the voxels in each data set.

For two volumetric data sets representing an object undergoing a simple rigid-body motion, the correspondence problem amounts to finding the six parameters of translation and rotation that best align the features in the two data sets. Motion of the heart throughout the cardiac cycle, however, is not rigid. The walls of the heart bend and stretch as an elastic body during the cardiac contraction. Therefore, in order to relate each voxel’s position in one data set to its corresponding voxel in another, a dense 3D vector map at the same resolution of the voxel size is required to describe the non-rigid deformation. This vector map is often called the motion field, or the flow field, which describes the non-rigid mapping. Though this map can adequately describe any non-rigid motion, it represents a tremendous increase over the rigid body case in the degrees of freedom allowed to describe the deformation.

The basic motion compensation scheme that is desired is depicted schematically in Figure 1-1. Here we see a single slice through a cardiac volume at end diastole and at end

systole. It is obvious from these images that the shape of the heart changes dramatically between these two time frames. We would like to select one of these volumes as a reference, and then warp the other data set, which we'll call the source volume, so that the shape of the heart in the deformed source volume matches the reference. The warping consists of two parts. First the motion field must be computed that defines the correspondence between each voxel in the source and reference volumes. Second, the motion vectors must be used to displace each voxel in the source volume and then recombine them in some manner to calculate the deformed volume. Once we have deformed this volume, the two volumes can then be summed together to form a composite volume with less motion blur and better contrast to noise characteristics.




Figure 1-1. Cardiac Motion Compensation Scheme

Cardiac motion is an elastic, non-rigid motion. To compensate for cardiac motion, separate images are acquired in a gated PET study from different phases of the cardiac cycle. To sum the different phases together, they must first be deformed to match a common reference phase. This establishes the correspondence between voxels representing matching portions of cardiac tissue. Once the correspondence is found and the shape of the heart matches in each deformed and reference volume, they may be summed together to form a composite data set with better contrast to noise properties.

The focus of this dissertation is the technique for performing the step that is italicized in the figure. That is, this dissertation will describe a technique used to estimate the motion field enabling a non-rigid deformation of a cardiac PET volume. The motion field will then be used to produce a motion compensated volume that may be subsequently added together as part of a deblurred composite volume. Though a gated cardiac sequence typically consists of many time frames, this dissertation will only focus on the combination of two of the frames to demonstrate the principle of non-rigid motion compensation. A source volume, typically the heart at end systole, will be deformed to match a reference volume, typically the heart at end diastole. The deformed source will be then summed with the reference to produce a composite volume with better contrast to noise characteristics. To limit the degrees of freedom in the motion field, the algorithm will make use of known physical properties of the cardiac tissue. This will prevent estimation of motions that would be physically impossible in the real world. The algorithm will be tested on phantoms modeled to mimic specific known characteristics of cardiac PET data, as well as on real data obtained from clinical cardiac PET studies.

1.2. Contribution of this Dissertation

Deformable motion algorithms that operate on density images have been in existence for over ten years now. Motivated primarily by the availability of high resolution 3D data sets acquired from medical imaging devices such as computed tomography (CT), magnetic resonance imaging (MRI) and PET, the deformable motion algorithms have become a necessary component of an analysis if one desires to compare elastic biological tissue imaged at different times, or to compare the image data from different subjects, which naturally has size and shape variations.

A common problem with the estimation of 3D deformable motion is how to properly deal with the huge dimensionality of the estimation problem. If one wants to warp the voxels in one volume to match another, there are nearly countless ways to do this and obtain an exact match. A simple example of this is seen frequently in the world of special effects for television and movies. Tigers have been smoothly “morphed” into automobiles, human faces have been transformed into “liquid metal,” and yet though we say that the voxels from one image have been warped to match the other, these warps do not repre-

sent a deformation that is physically possible. For transformations of the volumes acquired by medical imaging devices, we do not desire to use the freedom of possible motion fields used in the special effects community. It is known that the deforming images in cardiac PET actually correspond to physically deforming media. Therefore, use of a physical model to constrain the set of possible motion fields makes sense. The primary contribution of this dissertation is development of an appropriate physical model for the motion seen in cardiac PET data. The innovative aspects of the dissertation can be summarized as follows:

- Acknowledging that the objects in the PET scanner’s field of view do not have uniform material properties, we develop a piece-wise isotropic elastic material model. Past techniques have used material models, but the models were assumed to be isotropic.
- Deformations seen in cardiac tissue are large. Past techniques using ideas from material models have used what is called the infinitesimal motion assumption. This assumption may be appropriate for materials like steel, but they induce error for modeling the deformations seen in the heart. We will investigate the use of a *finite displacement* strain energy model that is more appropriate for describing large-scale deformations.
- The aspect of applying a deformable motion algorithm to decrease motion blur and improve overall image quality in cardiac PET is unique. Most past motion estimation techniques have been used to deform data sets so that corresponding pieces of tissue may just be compared. This work takes a second step to use the estimated flow field as a motion compensation technique.
- Because the motion field is being used to produce a warped image volume for the purpose of subsequently forming a composite data set, a *forward* warping scheme has been used to deform the source volume according to the motion field. This Lagrangian description of the motion field insures that every voxel in the source data set contributes to the deformed volume. Past techniques have used motion fields defined in an Eulerian sense, and have used *backward* warping schemes that had no such guarantee.

The work presented here is most closely related to 3D extensions of Horn’s optical flow algorithm, which made the assumption that corresponding voxels should have their

intensities conserved between two images in a motion sequence, and that the “optical flow field” describing the motion should be smoothly varying, just like the surfaces of the objects in the images are smoothly varying. Song and Leahy [82] were among the first to extend the optical flow algorithm to 3D, where they applied it to the motion estimation problem in cardiac ultrafast CT data sets. A separate body of work devoted to the warping of 3D brain data sets is also closely related to the approach explained in this dissertation. The brain warping algorithms make use of continuum mechanics and present material models to regularize the motion field estimation which are similar to those used in this work [7,14,92].

1.3. Summary of Results

The take-home message from this dissertation is that by more accurately modeling the elastic properties of the materials being imaged in cardiac PET, the motion estimation algorithm can be made more accurate. The estimated motion field can be used to compensate for the motion-induced blur in cardiac PET, and produce corrected images with less blur and better statistics.

It was found that use of a finite displacement strain energy model produced improvements over the strain energy model assuming infinitesimal motion. However, an even larger effect is the difference between the piece-wise isotropic model verses the completely isotropic model. In other words, it seems more important to model fairly incompressible cardiac tissue differently from the blood pool inside the heart than it is to model cardiac tissue elasticity differently than the elasticity of a material like steel. Also, the use of a forward sampling technique to guarantee the contribution of every voxel from a source volume during a deformation is an improvement to a backward sampling technique, though this becomes less important if the true motion field represents an incompressible deformation.

Finally, it should be noted that though the discussions in this dissertation have been targeted towards the modality of PET, most of the concepts described here are equally applicable to the deformation of data sets acquired with other imaging modalities.

1.4. Dissertation Outline

This dissertation is organized as follows. After the introduction found here, a comprehensive literature review is given in Chapter 2.

Chapter 3 explains the characteristics of typical gated cardiac PET data. The basic principles of PET acquisition and reconstruction are outlined here so that the reader can appreciate the logic that is used for assumptions about the image model. Examples of typical cardiac PET data are also shown in order to motivate the anisotropic strain model used in later chapters. It also should give the reader an understanding of some of the problems in cardiac PET imaging, and why this deformable motion algorithm is warranted.

Next, in Chapter 4, the general formulation for the deformable motion algorithm is stated. In this chapter are found most of the symbol definitions and other terms that are required for later discussions. The basic structure of the motion algorithm is outlined in terms of a similarity measure between the reference and deformed volumes and a regularization function for the motion field. Together, these criteria serve as a cost function describing a figure of merit for the overall motion estimation problem. A justification of the chosen similarity measure is next given. Then, past techniques used to regularize motion fields in optical flow-like problems are discussed, and we present a simple version of the linear elastic material model used in this work. A detailed discussion of the material model is reserved for Chapter 6. Two minimization methods are also presented in this chapter. The first of these is based upon the calculus of variations. It is used to minimize a version of the cost function based on the small motion assumption, which can be expressed in a quadratic form. The second minimization technique is a nonlinear successive over-relaxation technique. Though slower than the approach using the calculus of variations, this technique can be used to minimize more general cost functions.

Chapter 5 details some of the subtleties that can be encountered in the sampling process during a volume deformation. The motion vectors themselves are just one part of a volume deformation. The other component is the sampling technique which uses these vectors to displace each source voxel into a deformation volume. The chapter compares the forward sampling method here with techniques that have been used by other authors.

Chapter 6 gives a detailed explanation of the elastic material model used in the

deformable motion algorithm. Because the regularization constraint used by the algorithm is modeled from the characteristics of physically deforming elastic media, the relation of stress and strain in isotropic and nonisotropic media is explained, as well as how these quantities are used to calculate material strain energies appropriate for modeling cardiac tissue.

Next, Chapter 7 pulls together the concepts that have been explained in Chapters 4, 5 and 6, and summarizes them while presenting the final form of the deformable motion algorithm. It also gives some more of the implementation details, such as the multi-resolution approach that is used.

Results are presented in Chapter 8. A simple isotropic deformable phantom is first used to clearly demonstrate the behavior of the algorithm with respect to material model parameter choices, and how mismatches between the data and the chosen material parameters can affect the accuracy of the estimated motion. The Results chapter makes use of two anthropomorphic phantoms to evaluate the accuracy of the motion algorithm. Results using data from actual cardiac acquisitions are also presented.

A discussion of the results and plans for future work is found in Chapter 9.

Finally, an appendix is provided which gives details of the calculus of variations, and of the conjugate gradient minimization technique, both of which were used by the motion estimation algorithm. A basic explanation of these techniques is given, as well as an example of how they are used explicitly in this work.

Chapter 2

Related Work

2.1. Deformable Model Applications

Two general applications in the medical imaging community have generated considerable interest in deformable volume models. As the introduction in this document suggests, one application is the characterization of cardiac motion. For the PET application here, the motivation is to obtain a mapping which brings into correspondence those voxels from two different time frames representing the same cardiac tissue. Once the correspondence has been made, the image data may be added together to provide a composite image. Another motivation for the characterization of cardiac motion is the study of the motion itself, since the motion field can be used to study cardiac mechanics or indicate abnormal heart function based on wall motion, thickening or strain measurements [53,33,85,9].

The other general application which has driven research in deformable volume models is the warping of brain image data sets to match the shape of a standardized atlas. Recently available 3D imaging techniques, such as Magnetic Resonance Imaging, X-ray Computed Tomography and PET, provide high resolution *in vivo* data sets that allow routine visualization of the brain anatomy and physiological function. Because of natural size and shape variations between subjects, the brain data sets must be stretched or otherwise deformed to match the shape of a common atlas before comparisons can be made across subject populations.

Though the cardiac and neuroanatomical applications appear quite different, nearly

every technique employed to calculate the deformations for the two applications has been based on two general cost functions. One cost function is based on some similarity measure between the reference volume and the deformed source volume. The other cost measure is usually based on a smoothness measure of the motion field. A smoothness constraint is imposed because there are often numerous deformation fields that could produce a warped volume which is similar to or even exactly matching the reference volume; however, most of these deformations represent physically unrealizable warps. Therefore, an acceptable solution usually is a balance between the forces driving the similarity measure and the smoothness constraint. The similarity measure is the driving force that tries to make the source image deform until it matches the reference, and the smoothness constraint is the restraining force that keeps the motion field varying as if it represented the deformation of a physically deforming piece of elastic media, such as a piece of rubber undergoing some stress. The principle differences between the approaches lie in first what features are used to establish a similarity measure between the volumes; and second how the smoothness constraints are implemented.

Similarity measures are generally expressed in terms of either manually or automatically detected image features, or directly in terms of the voxels themselves. One of the simplest feature types that have been used as a similarity measure by many matching algorithms is a sparse set of paired control points. Typically, these points are specified by trained manual operators [30,11,24]. Alternately, the voxel data may be used directly as its own feature space. Common difference measures operating directly on the voxel data include the squared difference between individual voxels [103], uniformity of ratio [99], mutual information [57], or cross correlation [7]. Voxel-based measures have the advantage that they are usually automatically implemented, but they can be computationally intensive because they measure the difference at *every* voxel regardless of whether it is a relevant feature. As another approach, a dense similarity measure that can be quickly calculated may be used, such as the “daemon” technique of Thirion [90]. A compromise between manually specified points and the processing of every voxel may be obtained by preprocessing the volumes so that features such as edges, ridge lines [55,91,26], surface distance [64] or surface curvature [79,2,45] are used. These techniques attempt to extract the attributes of the data that are most useful for matching, although, in some cases it may

be difficult for a completely automatic preprocessing technique to extract features without destroying at least some potentially useful matching information.

Smoothness constraints can be classified somewhat by considering the total degrees of freedom allowed. At one end of the spectrum, the deformation field could be limited to six global rigid body motion parameters and three scale parameters. At the other end, the motion field could be defined by an independent vector at each voxel that links the position of that voxel to its corresponding location in the deformed volume. Smoothness constraints in this latter, high dimensional case are usually implemented as a function of the differential properties of this flow field.

2.2. Optical Flow Based Work

The algorithms in this dissertation are most closely related to works which find their roots in the two dimensional optical flow literature. In these bodies of work, a dense flow field is calculated by considering voxel intensity difference as a similarity measure. Horn [39] was one of the first to develop an optical flow algorithm that calculated the 2D vector field describing the motion between corresponding voxels in two images. The technique relied on an image matching constraint assuming corresponding pixels would have similar grey values in the two images, and a motion field smoothness constraint. The smoothness constraint was based on the assumption that objects in the images have continuous surfaces which induce a smoothly varying motion field.

Just as the work in this dissertation recognizes that elastic properties are not uniform in a cardiac PET volume, so too have other authors in the 2D optical flow world recognized that not all surfaces are continuous, and that motion fields describing discontinuous surfaces would violate a uniform smoothness constraint. To deal with discontinuous surfaces, some authors have used modified smoothness constraints that would not penalize discontinuities in motion fields that corresponded to surface boundaries. For example, Nagel [65] and Ghosal [35] investigated nonuniform smoothness constraints based on what they called an “oriented” smoothness criterion.

Similar works in the 3D deformation world parallel the formulations used by the 2D image algorithms. Fitzpatrick [27] demonstrated purely geometrical image transformations exist that have the same effect on the image as does the physical motion of the object

being imaged for 3D density images, such as those found in MRI and PET imagery. Song [82,83] used this information and a direct 3D extension of the Horn algorithm to calculate the motion field in ultrafast computed tomography (CT) images of the heart. Zhou [103] used a similar formulation to calculate larger deformations in CT images of deformed asphalt test structures. Both these techniques relied upon a voxel matching constraint as the driving force to the deformation. A simple smoothness constraint penalizing the square of the components of the flow field gradient was combined with an incompressibility constraint on the motion field to restrict the set of possible particle deformations. An analysis of the relative effectiveness of these constraints may be found in [36].

Bajcsy and others [7,21,34,60] were among the first to incorporate a more realistic elastic material model as a regularization constraint. Though they used their deformation technique to match 3D volumes of brains from different people, the brains appeared similar enough so that a material model of continuous media gave acceptable results. The linear elastic model assumed infinitesimal displacements, and since the deformations required to match different brain data sets can be quite large, it was a simplifying approximation to the actual material being imaged. Christensen [15,14] tried to overcome this problem by introducing a viscous fluid model capable of tracking large deformations. He used this technique to match largely differing brain data sets from different patients.

Other authors have used different similarity measures than voxel matching while retaining the high dimensionality of the motion vector field to describe the deformation. Davatzikos [22,23] made use of curvature matching and a deformable surface model to first register the cortical surface of a subject with a reference brain. This surface then served as the driving force to derive a 3D elastic map, using a similar linear elastic model as Bajcsy. Davis [24] also looked to physically based elastic constraint models, but used a sparse set of 20 manually identified landmarks to drive the deformation instead of voxel similarity measures. A similar approach was used by Thompson [92], who used a set of landmarks to fit a surface model called a Chen surface. The surface model was subsequently used to drive a smooth deformable sampling of the voxel data set. This technique produced deformations of brain data sets with impressive results.

2.3. Simple Regularization Constraints

Material model based algorithms appear as elegant solutions to deformable motion problems, however, they are computationally intensive because of the large dimensionality of the motion field solution set. Essentially, they make use of properties from real world elastic materials that enforce a smooth deformation of some 3D data set. This is necessary because without such constraints, nearly any arbitrary volume can be warped to match voxels in another volume in many ways. For warping brain data sets from different patients, it is known that the two brains are simply not the same piece of matter that has been warped. Rather, they are similar enough that a linear elastic or viscous fluid model can be used to approximate the required smooth deformation. However, there are a number of other, perhaps simpler, smoothness constraints that could similarly obtain a suitable deformation. For these reasons, a number of researchers have found success in the brain warping community using simpler models.

One of the simplest models is the nine-parameter conventional “Talairach” linear model [86], which allows three scaling parameters in addition to the six rigid-body registration parameters [30,70,17]. Simple global or piece-wise global models like these [61] or variants which allow affine [37,80] low-dimensional global polynomial or spline deformations [100,4,10,26] can be quite successful at matching data sets from different patients without gross anatomical abnormalities. These low-dimensional deformation models have been used extensively to register PET brain data sets for subsequent statistical analyses. Because relatively few parameters are required to specify the deformation, most of these techniques relied on manually specified landmarks. Global linear models have also been combined with local spline-based techniques with excellent results [62,93,32,31].

2.4. Miscellaneous Cardiac Motion Tracking Techniques

Because the shape of the heart is considerably simpler than the convoluted shape of the brain, many researchers have made use of this fact by including *a priori* parametrically described shape information in a deformable heart model. Cohen [16] used a deformable surface model inspired by the work of Terzopoulos’ active contour, or “snake” model [46,89] to first detect and then track deforming cardiac surfaces. Parametric surface models based on superquadrics [8,13,88,92], spherical harmonics [56], polar transforma-

tions [25], or probabilistic models [84] have all been used to assist in the segmentation and tracking of the cardiac surface. Finite element models of varying detail have also been used to describe the motion of the heart [71]. Some go as far as mapping the cardiac muscle fiber direction and use sophisticated material stress/strain models to predict the motion [59,95,42,66,20,43]. These models, though computationally intensive, can model the motion of the heart extremely well.

Finally, a considerable amount of work in deformable volume models has been the result of imaging techniques specifically designed to study cardiac motion. MRI tagging is one such technique. With this type of imaging, a spatially varying pattern, or tag, is induced in the data at the start of each cardiac cycle. Motion of the tissue after each tag application is revealed by tracking the tag pattern [102,6]. Several methods have been developed to track these tags [73,101,1,74,67]. Another MRI technique called phase contrast MRI, which produces images of cardiac tissue velocity [69], has recently emerged as a technique to study heart motion, and has spawned a variety of deformable modeling techniques which attempt to characterize the cardiac motion [18,58]. Techniques such as these differ considerably in the feature mapping component of the deformable model algorithm, but they share much in common with the techniques used to constrain the solution set of allowable motions.

2.5. Perspectives on Related Work

An optical flow-based technique is chosen for the application of motion compensation in cardiac PET for a number of reasons. The technique naturally results in a dense estimation of motion vectors that can be used to describe the motion of every voxel in a source data set without need for interpolation. Use of the voxel data as a similarity measure instead of manually defined or automatically detected image features reduces the requirement for image preprocessing, and avoids the potential problem of a feature detection algorithm throwing away useful data too quickly.

In terms of computational requirements, the estimation of a dense motion vector field at the same resolution of the image data set makes this a more expensive technique than motion estimation algorithms based on a simplified parametric motion model, such a linear, affine or polynomial warp of the data. On the other hand, the dense motion field is

much more capable of describing general deformations, and investigations of more complex motion types are possible. At the other end of the spectrum, because the algorithm described in this thesis requires little image preprocessing and object recognition, it is considerably simpler than the detailed finite element models used by Nielsen [66] and others to describe cardiac deformations.

Chapter 3

Imaging Model

Since this work exploits the properties of cardiac data imaged through PET in an effort to correct for motion, it is necessary to understand at least the fundamental principals involved in obtaining a gated cardiac PET reconstruction. Positron Emission Tomography (PET) is a technique used to image radioactive tracers. The tracers are designed so that their chemical properties are nearly identical to naturally occurring compounds existing in the body. They can therefore be used to tag some biochemical process of interest, such as metabolism, ligand-receptor interactions or blood flow. A commonly used tracer in cardiac studies is ^{18}F -fluorodeoxyglucose (FDG). This compound is a glucose analog that indicates metabolic uptake of glucose by myocardial tissue in walls of the heart. In a typical FDG cardiac study, the tracer is administered intravenously a short time before the study is begun. As the body metabolizes glucose, the FDG is taken up by the tissue and remains trapped there, while the remaining compound in the bloodstream is washed out relatively quickly. PET data are usually acquired after the tracer has cleared the bloodstream so that the only radioisotope in the scanner's field of view is that which is "trapped" in the tissue. In the case of cardiac imaging, the left ventricle uses proportionately much more glucose than other nearby tissue does, so that this object is the prominent feature in a reconstructed image volume. Seen in Figure 3-1, the left ventricle appears in a transverse slice through the torso as a cup-shaped object far brighter than any adjacent tissue, including the right ventricle or other portions of the heart.




Figure 3-1. Typical FDG Cardiac PET Image

Transverse slice through the body shows that the left ventricle appears as a cup shaped object that is the principal feature in the PET image. Activity in the blood pool, right ventricle and lungs is relatively low.

3.1. Tomography - PET Image Reconstruction

All radioisotopes used in PET decay by emitting positrons. These positrons travel a very short distance (typically less than 5 mm) from the decay site and then annihilate as they meet an electron, forming two 511-keV gamma rays that are emitted simultaneously at directions nearly 180 degrees with respect to each other. The gamma rays are detected by rings of scintillation crystals in the scanner, and are tagged as an annihilation pair by using extremely fast coincidence electronics. Each detected pair therefore indicates that a positron annihilation event has occurred somewhere along the line joining the two detection crystals. The direction of the emitted gamma rays is completely random for each annihilation, but by recording enough of these events, a complete projection data set may be obtained that measures the event density at all angular and linear positions. This projection data set, often called a sinogram, indicates the integral of all tracer activity along each line of response. It can be reconstructed by a number of tomographic techniques into a 3D image volume. Details about various reconstruction techniques are beyond the scope of this dissertation. See [40] for an in-depth description of many reconstruction techniques. The facts that are relevant to this work can be summarized as follows:

- Reconstructed PET voxel intensity represents the estimated concentration of radioactive tracer present within that voxel.
- The noise characteristics of the reconstruction are largely a function of the total tomograph events collected during an acquisition. More events result in a better statistical estimate of the true tracer concentration in each voxel.
- Tomograph events can be stored in different sinograms at different times, then reconstructed into separate volumes to indicate the tracer concentration at different time points. However, because the total number of events has been distributed into more than one projection set, the resulting reconstructions appear noisier.

3.2. Gated PET Acquisitions

In gated cardiac acquisitions, the cardiac cycle is monitored with an electrocardiogram (ECG) so that the R-wave at the start of each cycle is detected. Using this time point as a reference, tomograph events are then directed into different storage locations, dividing the cardiac cycle into some 5-15 time segments, called gates, each typically 50-100 msec in length. Because a single 100 msec time segment would not provide nearly enough events for an adequate image reconstruction, the events from a large number of cardiac cycles are summed over the course of a 5 - 30 minute scan. In its simplest form, cardiac gating is carried out by using the same time segments with respect to the R-wave for each gate during every cardiac cycle. This simple scheme is usually adequate, though in some cases where the heart beat is extremely irregular, the gating may be improved by using other approaches [63,52]. The overall goal of the gating is to provide a sequence of projection data sets where each represent the shape of the heart at a specific phase of the cardiac cycle. Hence, the reconstructed volumes effectively capture the heart motion as a series of freeze frames.

PET data used as test data sets in this work were acquired on the CTI/Siemens ECAT EXACT HR tomograph. The scanner has an in-plane peak resolution of 3.6 mm full width at half maximum (FWHM) and a peak axial resolution of 4.0 mm [97]. Rather than binning tomograph events directly into projection arrays, a list mode acquisition scheme was used which stored events individually to a disk in real time. This allowed a retrospective




Figure 3-2. Ungated Cardiac Data

Three orthogonal views of an ungated cardiac study. Reconstructed emission volume seen in (a) and the transmission correction map (b) used to correct the emission data for the effects of positron attenuation. The reconstruction in (c) shows the same 3 slices as in (a), but using a voxel size and reconstruction offsets zoomed in on the left ventricle. Displays of 3D data sets will often be displayed in this format. To aid in the identification of familiar anatomical structures, orientations of the images with respect to body position are given. The intersection of the other image planes with each image is also shown.

gating scheme, where the data could be replayed at will using different cardiac gating or other preprocessing schemes. For details see [49,41,50].

An example of three orthogonal slices through an ungated cardiac study acquired with this system is seen in Figure 3-2. This image volume may be compared with a gated study in Figure 3-3 that was created from the same tomograph data using 100 msec gates and




Figure 3-3. Gated Cardiac Data

Three orthogonal views of the first four time frames in a gated cardiac study. Gate duration is 100 msec. Though there is less blur in the images than in the ungated case, the images are considerably noisier.

reconstructed using identical smoothing filters as were used in the ungated case. As expected, the ungated sequence is considerably noisier. However, one can see even through the noise in the gated sequence the motion that takes place between diastole and systole. Typically, this motion is usually in the range of 1-13 mm [77,67,28,29], which is also in the range of the thickness of the left ventricle wall [53]. So even though conventional PET scanners can resolve objects smaller than the thickness of the heart walls, the motion in ungated data sets severely limits the effective spatial resolution.

3.3. Imaging Model Implications

As one would expect, the tissue of the heart bends and stretches during the cardiac cycle in accordance with the laws governing physical elastic media. Given that cardiac imaging takes place after the radiotracer has left the bloodstream, nearly all the activity within the field of view of the tomograph is trapped in the tissue, and should therefore be approximately conserved. Motion seen in the voxels between two image volumes must therefore be in agreement with the motion of the underlying elastic tissue. Because the myocardial tissue is approximately incompressible [59,19], it should be expected that the motion field within the cardiac field should be divergence-free. Note however, that though the tissue being imaged in a cardiac study is a continuum, the elastic properties of adjacent regions in the image would be considerably different, such as the blood pool within the ventricles of the heart and the lung air space. Indeed, the blood pool is not only a viscous fluid with different material properties from the myocardial tissue, but it also is an object that moves in and out of the field of view during the cardiac cycle. Therefore, though elastic constraints should be enforced on the motion field for voxels representing the same tissue type, they do not necessarily need to hold across tissue boundaries.

These characteristics differ from many other applications of 3D deformations, for example, in the matching of MRI brain data sets to a common brain atlas. In this latter case, the two objects being matched are not physically the same piece of tissue. Therefore one would not necessarily expect image density to be conserved, nor would motion necessarily have to conform to elastic material models.

Chapter 4

General Formulation

4.1. Deformation Framework

The deformable modeling algorithm described here employs a densely defined vector field to non-rigidly displace voxels from a source volume, thus forming a deformed volume. The goal of the algorithm is to calculate the motion vector field such that the deformed volume closely matches a third data set called the reference volume. A cost function that naturally follows from this goal is an image matching constraint measuring in some way the difference between the reference volume and the warped source volume. The dimensionality of the motion vector field is extremely high – a three-component vector is defined to describe the motion at each voxel. Due to this, there are frequently numerous motion fields that could be used to warp the source volume and equally satisfy the image matching constraint. Therefore, an additional cost function is usually included which enforces some measure of motion field smoothness to regularize the solution. To make a physical analogy, the image matching constraint can be thought of as a stretching and bending force that tries to pull voxels in the source volume to their corresponding location in the reference volume, and the smoothness constraint can be viewed as a set of tiny springs attached between adjacent voxels which resist any deformation due to the image matching force. Indeed, this physical analogy is just what the current formulation attempts to model. Factors weighting the cost of motion field smoothness violations are chosen so that the elastic restraining force is similar to the physical restraining forces seen in the car-

diac tissue being imaged. The elastic restraining force of the smoothness constraint prevents any deformations that are physically unrealizable, thus considerably reducing the domain of allowable motion fields. Typically, there is noise present in the source and reference image volumes. The restraining forces are also used to prevent false correspondence between two uncorrelated intensity peaks due purely to noise.

The motion estimation framework is described as follows. Define two 3D density fields, $f_1(\mathbf{r})$ and $f_2(\mathbf{r})$, where $\mathbf{r} = (x, y, z)$ represents the voxel index in a discrete domain, $\mathbf{r} \in \{[1, N_1], [1, N_2], [1, N_3]\}$ and N_1, N_2, N_3 are the dimensions of the image density fields. We will call $f_1(\mathbf{r})$ the source volume, and $f_2(\mathbf{r})$ the reference volume. A Lagrangian motion field is defined as,

$$\mathbf{m}(x, y, z) = (u(x, y, z), v(x, y, z), w(x, y, z)) \quad (4-1)$$

and the deformed volume of f_1 is defined as $\hat{f}(\mathbf{r} + \mathbf{m}(\mathbf{r})) = f_1(\mathbf{r})$. We will call $\hat{f}(\mathbf{r})$ the deformed or the target volume. As depicted in Figure 4-1, each motion vector, $\mathbf{m}(\mathbf{r})$,




Figure 4-1. Voxel Correspondence

A Lagrangian motion vector field is defined that expresses the mapping between corresponding voxels in the source and reference volumes.

defines the displacement necessary to bring a voxel from the source volume into correspondence with the voxel representing the same tissue in the reference volume. In a continuous sense, the motion field must therefore be one-to-one and onto, thus defining a homeomorphic transformation. Assuming f_1 and f_2 are measurements of a conserved medium, the goal of the motion estimation technique is to find the physically plausible motion field such that $\hat{f}(\mathbf{r}) \approx f_2(\mathbf{r})$.

With these definitions, we can express generalized cost functions for the image matching and smoothness constraints at each voxel location, \mathbf{r} , as follows:

$$e_I(\mathbf{r}) = g_I(\hat{f}(\mathbf{r}), f_2(\mathbf{r})) \quad (4-2)$$

and

$$e_S(\mathbf{r}) = g_S(u_x, u_y, u_z, v_x, v_y, v_z, w_x, w_y, w_z, u_{xx}, u_{xy}, \dots, w_{zz}) \quad (4-3)$$

where partial derivatives of the motion field are written using the notation: $u_x \equiv \partial u / \partial x$, $u_{xy} \equiv \partial^2 u / \partial x \partial y$, g_I and g_S represent generic functions chosen for the image matching and regularization constraints. To simplify notation, the \mathbf{r} argument has been dropped in the partial derivative terms. So the image matching constraint is a function of voxel differences between the reference and deformed source volumes, and the smoothness constraint is a function of the first and second order differential properties of the motion field. The overall minimization problem is to find a motion field consistent with elastic material properties that best match the deformed and reference via a minimization of:

$$E_{tot} = \sum_{\mathbf{r}} [e_I(\mathbf{r}) + e_S(\mathbf{r})] \quad (4-4)$$

4.2. Similarity Measures

Ideally, a good similarity measure for the cardiac PET application should be efficient to calculate, should be valid for possible long-range displacements, and should be relatively robust in the sense that it should provide a good measure of the closeness between a reference and a deformed volume over a wide range of noise conditions. Also, because a

motion vector must be estimated at every voxel, the similarity measure must provide a measure of image difference at every voxel in the deformed data set. This latter requirement precludes many manually driven comparisons of a sparse feature set.

Early optical flow methods were based on the assumption that the intensity of a pixel (or voxel in 3D) is proportional to a quantity that is conserved as the underlying medium moves. For example, in the case of 2D imaging, the quantity was the reflectance at a given point on a surface. For 3D materials, the conserved quantity could be the proton density for MRI data, or the tracer concentration in PET data. We can view the voxel intensity representing this conserved quantity in a deforming volume as the time varying function, $f(\mathbf{r}, t)$. Assuming also that the underlying medium is incompressible leads to the following similarity measure in differential form:

$$f_x u + f_y v + f_z w + f_t = 0 \quad (4-5)$$

This is the 3D equivalent to the brightness constraint of Horn [39]. It has been used as a similarity measure for 3D data sets by Song [82] and others. One can easily see in this equation the ill-posed nature of the problem estimating the motion field. There are three unknown variables and only one equation. Still, by combining this equation with simple smoothness constraints, the differential form of the similarity measure has the advantage that it lends itself naturally to well-posed linear regression minimization techniques. For the present application of cardiac PET, however, the measure is inadequate because inter-frame displacements can be considerably larger than one voxel, and hence Equation 4-5 is not applicable.

For long distance displacements, it is natural to look at similarity measures between voxels in the reference and in the deformed volume. A rather brute force method of implementing this, which has been used by Bajcsy [7] and others, is a normalized cross correlation similarity measure:

$$e_I(\mathbf{r}) = \frac{\sum_{\tilde{\mathbf{r}} \in R_r} f_2(\tilde{\mathbf{r}}) \hat{f}(\tilde{\mathbf{r}})}{\sqrt{\sum_{\tilde{\mathbf{r}} \in R_r} (f_2(\tilde{\mathbf{r}}))^2} \sqrt{\sum_{\tilde{\mathbf{r}} \in R_r} (\hat{f}(\tilde{\mathbf{r}}))^2}} \quad (4-6)$$

This technique is useful for measuring the difference between two image volumes containing voxels representing the same type of material. Indeed, it would be a useful measure for comparing different gates from a cardiac PET sequence; however, it is computationally expensive. If simpler measures are available that are still reliable, it would be better to use them.

Numerous other similarity measures have been reported that are applicable to deformations with long-distance displacements, and have varying degrees of robustness and requirements for user interaction and computer processing. For example, curvature matching techniques have been used by some investigators. The geometry of the left ventricle can be described by two well-behaved surfaces, the endocardium is the inner surface, and the epicardium is the outer surface. By modeling the stretching of the heart as a conformal stretch, measures such as Gaussian curvature may be used to gauge surface point correspondence [2,45]. Surface-based measures like these can be successful provided that the surfaces can be estimated accurately. However, this is not a trivial task in PET data because of the noise typically present. The fact that curvature relies on high order partial derivatives of the parameterized surface makes the technique vulnerable to error in high noise settings. Also, like most other feature extraction techniques, the reliability of the similarity measure is extremely dependent on the reliability of the technique to extract similar features in the two data sets. That is, though the feature extraction simplifies subsequent comparison calculations, some important information may have been lost during the feature extraction process. These techniques using a sparse set of extracted features also need to provide a method for interpolating the difference measure at the locations of each voxel.

One obvious simple measure that does not require any feature extraction and is valid for large interframe displacements is the squared difference measure between single voxels:

$$e_I(\mathbf{r}) = \gamma_I(f_2(\mathbf{r}) - \hat{f}(\mathbf{r}))^2 \quad (4-7)$$

where γ_I is a global scalar used to alter the balance between the image matching and regularization terms. We choose this measure because it is computationally efficient, it does

not throw away potentially useful image information, and it does not require any operator intervention. Also, because it provides a similarity measure at every voxel instead of a sparse set of feature points, techniques are not necessary to interpolate values between features.

4.3. Regularization Constraints

As Horn noted in his 2D work on optical flow, the flow field problem based only on the brightness constraint is ill-posed. It therefore requires an extra regularization constraint to obtain a unique solution. The regularization constraint he proposed penalized departures from flow field smoothness by weighting the square of the first partial derivatives of the 2D flow field as follows:

$$e_S(\mathbf{r}) = u_x^2 + u_y^2 + v_x^2 + v_y^2 \quad (4-8)$$

A 3D generalization expressing all squared components of the flow field gradient in this simple smoothness constraint is

$$e_S(\mathbf{r}) = u_x^2 + u_y^2 + u_z^2 + v_x^2 + v_y^2 + v_z^2 + w_x^2 + w_y^2 + w_z^2 \quad (4-9)$$

Song and Leahy [82,83] and later Zhou [103] combined this constraint with a divergence-free constraint:

$$e_D(\mathbf{r}) = \lambda_D(u_x + v_y + w_z)^2 \quad (4-10)$$

which in essence modeled the underlying material as conserved, incompressible medium if we require $e_D(\mathbf{r})$ to equal zero at every voxel. Rather than making the divergence penalty a hard constraint limiting the flow field to zero divergence, both Song and Zhou chose to weight the constraint as an error term which is minimized along with the image matching and smoothness criteria. Combining Equation 4-5 with Equation 4-9 and Equation 4-10 in this manner results in a well-posed problem that can be solved via standard minimization techniques.

In terms of a material model for the medium being imaged, the only assumption made using the smoothness and divergence constraints discussed thus far is that the medium is a conserved quantity which resists compression and deforms smoothly. Other authors have

used more sophisticated constraints based on continuum theory and material models. A considerable body of work modeling deformable curves and surfaces known as *snakes* was based on two types of material constraints [46,16]. For a parameterized two-dimensional snake surface embedded in a three-dimensional volume, $\mathbf{v}(s, r) = (v_1(s, r), v_2(s, r), v_3(s, r))$, an internal energy function was defined by:

$$E_{int} = \omega_{10} \left\| \frac{\partial \mathbf{v}}{\partial s} \right\|^2 + \omega_{01} \left\| \frac{\partial \mathbf{v}}{\partial r} \right\|^2 + \omega_{20} \left\| \frac{\partial^2 \mathbf{v}}{\partial s^2} \right\|^2 + \omega_{11} \left\| \frac{\partial^2 \mathbf{v}}{\partial s \partial r} \right\|^2 + \omega_{02} \left\| \frac{\partial^2 \mathbf{v}}{\partial r^2} \right\|^2 \quad (4-11)$$

The first order derivative terms in this equation are called the membrane terms, for they make the surface resistant to stretching, much like an elastic string in one dimension, or a membrane in two dimensions. The second order derivative terms make the model resistant to bending forces, and can thus be interpreted as a rod in one dimension or a thin plate model in two dimensions. The individual weighting terms, ω_{ij} , are used to control the elastic properties of the model. Obviously, increasing the value of the first order derivative weighting terms make the model more resistant to stretching, and likewise, increasing the second order terms make the model more resistant to bending.

These two physically inspired restraining forces have appeared together or independently in a number of other applications, both for modeling parametric surfaces and 3D solid models. If we set the snake weighting factors to one and develop the restraining forces for an isotropic solid, we arrive at the equation in three dimensions for the membrane term [3]:

$$e_M(\mathbf{r}) = u_x^2 + u_y^2 + u_z^2 + v_x^2 + v_y^2 + v_z^2 + w_x^2 + w_y^2 + w_z^2 \quad (4-12)$$

Because the partial differential equations expressing the equilibrium condition for a body undergoing stress using this model result in a Laplacian operator, this model is also known as the Laplacian model. We also see that the snake membrane term is exactly the same as Horn's and Song's smoothness criteria.

The thin-plate energy function has also been used alone as a restraining force by a number of authors [10,24,78]. Generalized for a volume in three dimensions, this model, also known as the biharmonic, or the thin plate spline model can be expressed as:

$$e_T(\mathbf{r}) = (u_{xx}^2 + u_{yy}^2 + u_{zz}^2) + 2(u_{xy}^2 + u_{yz}^2 + u_{xz}^2) + \quad (4-13)$$

$$(v_{xx}^2 + v_{yy}^2 + v_{zz}^2) + 2(v_{xy}^2 + v_{yz}^2 + v_{xz}^2) +$$

$$(w_{xx}^2 + w_{yy}^2 + w_{zz}^2) + 2(w_{xy}^2 + w_{yz}^2 + w_{xz}^2)$$

The three dimensional generalizations of the membrane model and the thin plate spline certainly restrain a volumetric model from deforming, but a physical interpretation of them is not always readily intuitive since they are based on one and two-dimensional physical models. Many papers modeling solid object deformation have been written using a more intuitive formulation based on continuum mechanics called the linear elastic model [12,7,60,34,24,22]:

$$e_S(\mathbf{r}) = \frac{\lambda}{2}(u_x + v_y + w_z)^2 + \mu(u_x^2 + v_y^2 + w_z^2) \quad (4-14)$$

$$+ \frac{\mu}{2}(u_y^2 + u_z^2 + v_x^2 + v_z^2 + w_x^2 + w_y^2 + 2u_y v_x + 2u_z w_x + 2v_z w_y)$$

where λ and μ are elastic weighting terms called the Lamé constants. These can be written in terms of the more intuitive constants, E , called the Young's elasticity modulus and ν , called the Poisson ratio:

$$E = \frac{\mu(3\lambda + 2\mu)}{(\lambda + \mu)} \text{ and } \nu = \frac{\lambda}{2(\lambda + \mu)} \quad (4-15)$$

E relates the tension of the object to its stretch in the same direction, and ν is the ratio between lateral contraction and axial extension [5]. Comparing with the simpler constraints of Equation 4-8 and Equation 4-9, it can be seen that the λ term in the equation penalizes non-zero divergence and the μ term penalizes sharp discontinuities in the flow field. For highly incompressible fields, the Poisson ratio approaches a maximum of 0.5, which yields a divergence term, λ , that approaches infinity. The linear elastic model has often been expressed in the form of the equilibrium partial differential equations for a homogeneous, isotropic body undergoing stress:

$$\mu \nabla^2 \mathbf{m} + (\mu + \lambda) \nabla(\nabla \bullet \mathbf{m}) = \mathbf{F} \quad (4-16)$$

where \mathbf{F} is a distributed external force applied to the body, $\nabla \bullet \mathbf{m}$ is the divergence of \mathbf{m} , $\nabla^2 \mathbf{m}$ is the Laplacian of \mathbf{m} , and ∇ is the gradient operator. This expression is known as the Navier equation.

For three-dimensional deformations of biological tissue, such as MRI data sets of breast tissue, the linear elastic model has been found to be slightly more accurate than the thin plate spline [24], though perhaps this is at the cost of extra processing time and the requirement for specifying the Poisson ratio and the elasticity modulus. Choice of the elasticity parameters varies according to the application. For modeling breast tissue, Davis chose values of the Poisson ratio to be 0.01, 0.25 and 0.49. Of these values, the ratio of 0.25 produced the best matching, indicating that modeling the tissue as somewhat compressible was optimal [24]. For other applications, such as the matching of brains from different subjects, there is no intuitive reason to expect one value of the Poisson ratio to work better than another, since volumes being deformed represent two entirely different physical objects. Perhaps for this reason, Bajcsy set the Poisson ratio to zero to allow greater freedom in the elastic matching process [7].

One shortcoming in all the models discussed thus far for modeling biological deformations is that they assume infinitesimal deformations. This may be appropriate for materials like steel, but they do not necessarily apply for the deformations that take place in the heart, or for the matching of two different brain data sets. The models using an infinitesimal displacement assumption will still impose a penalty for a volume deforming with large deformations, and enforce the same topological properties between a deformed and a reference volume, but an intuitive interpretation of the elasticity parameters may be misleading. An attempt to more accurately model large scale deformations has been carried out by Christensen [15, 14]. In his work, which was applied to the matching of different MRI brain data sets, he modeled the volume as a viscous fluid. The partial differential equation modeling this behavior is as follows:

$$\alpha \nabla^2 \mathbf{v}(\mathbf{r}, t) + (\alpha + \beta) \nabla (\nabla \bullet \mathbf{v}(\mathbf{r}, t)) = \mathbf{F} \quad (4-17)$$

where the deformation *velocity* is given by

$$\mathbf{v}(\mathbf{r}, t) = \frac{d}{dt}\mathbf{m}(\mathbf{r}, t) = \frac{\partial \mathbf{m}}{\partial t} + \mathbf{v}(\mathbf{r}, t)\frac{\partial \mathbf{m}}{\partial x} + \mathbf{v}(\mathbf{r}, t)\frac{\partial \mathbf{m}}{\partial y} + \mathbf{v}(\mathbf{r}, t)\frac{\partial \mathbf{m}}{\partial z} \quad (4-18)$$

where α and β are viscosity constants, and where \mathbf{F} is a driving force that decreases as the deformed volume approaches the shape of the reference. Note that Equation 4-17 bears more than just a resemblance to the Navier equation describing the linear elastic model, however, by solving this equation successively at a large number of time points, the deformed volume could be made to track large displacements.

The viscous fluid approach was found to be very successful for matching largely differing brain data sets without inducing a topology change that would be comparable to a discontinuous tear in the deformed volume. However, it is probably not the most appropriate model for cardiac tissue. Though the walls of the heart can deform with displacements that are larger than an infinitesimal formulation would model, they are certainly less compliant than a viscous fluid. Christensen's approach allowed for sinks and sources in the model; a characteristic that would violate the conservation of activity assumption for cardiac PET data. Also, the viscous fluid method is extremely computationally expensive. For these reasons, we will begin by using the linear elastic model for the work in this thesis. Anisotropic and large displacement variations based on this model will be discussed in Chapter 6.

4.4. Criteria Minimization

4.4.1. Euler-Lagrange Approach

Recall that in order to find the motion field that best deforms a volume to match a reference, we are minimizing the total energy as defined in Equation 4-4, which is the sum over all voxels of the image matching and regularization criteria. Note, that when using a Lagrangian motion description, the image matching terms are evaluated at the “forward warped” location, $\mathbf{r} + \mathbf{m}(\mathbf{r})$. Details on the calculation of the deformed volume, $\hat{f}(\mathbf{r})$, given the motion field, \mathbf{m} , are found in Chapter 5. When using the voxel squared difference image matching criterion and an isotropic linearly elastic material model, this energy function is:

$$\begin{aligned}
E_{tot}(\mathbf{m}(\mathbf{r})) = \sum_{\mathbf{r}} & \left\{ \gamma_I [f_2(\mathbf{r} + \mathbf{m}(\mathbf{r})) - \hat{f}(\mathbf{r} + \mathbf{m}(\mathbf{r}))]^2 + \right. \\
& \frac{\lambda}{2} (u_x + v_y + w_z)^2 + \mu (u_x^2 + v_y^2 + w_z^2) + \\
& \left. \frac{\mu}{2} (u_y^2 + u_z^2 + v_x^2 + v_z^2 + w_x^2 + w_y^2 + 2u_y v_x + 2u_z w_x + 2v_z w_y) \right\}
\end{aligned} \tag{4-19}$$

One approach to this optimization is to invoke a conjugate gradient minimization technique similar to the method proposed by Zhou [103]. The non-linear process of calculating the best deformed volume matching the reference is linearized using a Taylor series approximation. Assuming the true motion field is \mathbf{m} , and the current estimate is $\tilde{\mathbf{m}}$, then a Taylor series approximation of $\hat{f}(\mathbf{r} + \mathbf{m}(\mathbf{r}))$ can be expressed in terms of a delta flow field, $(\delta u, \delta v, \delta w) = \delta \mathbf{m} = \tilde{\mathbf{m}} - \mathbf{m}$ as:

$$\hat{f}(\mathbf{r} + \mathbf{m}(\mathbf{r})) = \hat{f}(\mathbf{r} + \tilde{\mathbf{m}}) - \nabla \hat{f}(\mathbf{r} + \tilde{\mathbf{m}}) \bullet \delta \mathbf{m} \tag{4-20}$$

Substituting the expression, $\tilde{\mathbf{m}} - \delta \mathbf{m}$, for \mathbf{m} in the constraint equations results in

$$\begin{aligned}
E_{tot}(\mathbf{m}) = \sum_{\mathbf{r}} & \left\{ \gamma_I [f_2(\mathbf{r} + \mathbf{m}) - \hat{f}(\mathbf{r} + \tilde{\mathbf{m}}) + \nabla \hat{f}(\mathbf{r} + \tilde{\mathbf{m}}) \bullet \delta \mathbf{m}]^2 + \right. \\
& \left[\frac{\lambda}{2} (\tilde{u}_x - \delta u_x + \tilde{v}_y - \delta v_y + \tilde{w}_z - \delta w_z)^2 + \right. \\
& \mu ((\tilde{u}_x - \delta u_x)^2 + (\tilde{v}_y - \delta v_y)^2 + (\tilde{w}_z - \delta w_z)^2) + \\
& \frac{\mu}{2} ((\tilde{u}_y - \delta u_y)^2 + (\tilde{u}_z - \delta u_z)^2 + (\tilde{v}_x - \delta v_x)^2 + (\tilde{v}_z - \delta v_z)^2 + (w_x - \delta w_x)^2 \\
& + (\tilde{w}_y - \delta w_y)^2 + 2(\tilde{u}_y - \delta u_y)(\tilde{v}_x - \delta v_x) + 2(\tilde{u}_z - \delta u_z)(w_x - \delta w_x) \\
& \left. \left. + 2(\tilde{v}_z - \delta v_z)(\tilde{w}_y - \delta w_y) \right) \right] \left. \right\}
\end{aligned} \tag{4-21}$$

Equation 4-21 represents a quadratic functional in $\delta \mathbf{m}$ that can be minimized using the calculus of variations [38] (See Appendix A). The resulting Euler-Lagrange equations that can be derived for each component of the motion field are as follows:

$$\begin{aligned} \gamma_I(f_2 - \hat{f} + \nabla \hat{f} \bullet \delta \mathbf{m}) \hat{f}_x = & \lambda(\delta u_{xx} + \delta v_{xy} + \delta w_{xz} - \tilde{u}_{xx} - \tilde{v}_{xy} - \tilde{w}_{xz}) \\ & + \mu(2\delta u_{xx} + \delta u_{yy} + \delta u_{zz} + \delta v_{xy} + \delta w_{xz} - 2\tilde{u}_{xx} - \tilde{u}_{yy} - \tilde{u}_{zz} - v_{xy} - w_{xz}) \end{aligned} \quad (4-22)$$

$$\begin{aligned} \gamma_I(f_2 - \hat{f} + \nabla \hat{f} \bullet \delta \mathbf{m}) \hat{f}_y = & \lambda(\delta u_{xy} + \delta v_{yy} + \delta w_{yz} - \tilde{u}_{xy} - \tilde{v}_{yy} - \tilde{w}_{yz}) \\ & + \mu(\delta v_{xx} + 2\delta v_{yy} + \delta v_{zz} + \delta u_{xy} + \delta w_{yz} - \tilde{v}_{xx} - 2\tilde{v}_{yy} - \tilde{v}_{zz} - u_{xy} - w_{yz}) \end{aligned} \quad (4-23)$$

$$\begin{aligned} \gamma_I(f_2 - \hat{f} + \nabla \hat{f} \bullet \delta \mathbf{m}) \hat{f}_z = & \lambda(\delta u_{xz} + \delta v_{yz} + \delta w_{zz} - \tilde{u}_{xz} - \tilde{v}_{yz} - \tilde{w}_{zz}) \\ & + \mu(\delta w_{xx} + \delta w_{yy} + 2\delta w_{zz} + \delta u_{xz} + \delta v_{yz} - \tilde{w}_{xx} - \tilde{w}_{yy} - 2\tilde{w}_{zz} - u_{xz} - v_{yz}) \end{aligned} \quad (4-24)$$

These equations are linear in $\delta \mathbf{m}$ and are solved via standard finite differencing techniques and a conjugate gradient algorithm. Solution of the overall problem is thus carried out in a two-step process involving two loops. In the outer loop, the estimate of the motion field, $\tilde{\mathbf{m}}$ is used to calculate the current deformation volume, $\hat{f}(\mathbf{r})$. Then, in an inner loop, a conjugate gradient algorithm is used to find the best $\delta \mathbf{m}$ satisfying Equation 4-21. This delta motion field is added to the current total flow field and the procedure is repeated. For the results presented in this dissertation, ten to fifteen iterations of this outer loop were typically required to reach an overall solution. Each conjugate gradient step usually converged quickly, and also required some ten to twenty iterations.

4.4.2. Nonlinear Over-Relaxation Method

An alternate minimization technique that is a bit more straightforward than the previous approach is based on nonlinear successive over-relaxation (NLOR) [81]. Rather than using a Taylor approximation to obtain a linear formulation, the NLOR technique computes derivatives of the energy function directly to minimize any linear or nonlinear system of the form

$$g_k(x_1, x_2, \dots, x_M) = 0, \quad k = 1, \dots, M \quad (4-25)$$

The minimization technique can therefore be used to investigate formulations of this deformation problem that include nonlinear functions.

For each of the $M = 3N_1N_2N_3$ parameters in the motion vector field, we can define the derivative of the energy function with respect to that parameter as:

$$g_k(u_1, v_1, w_1, \dots, w_{N_x N_y N_z}) = g_k(x_1, x_2, \dots, x_M) = \frac{\partial E_{tot}}{\partial x_k} \quad (4-26)$$

A NLOR iteration is defined as:

$$x_k^{n+1} = x_k^n + \omega \frac{g_k(x_1^{n+1}, \dots, x_{k-1}^{n+1}, x_k^n, \dots, x_M^n)}{\frac{\partial}{\partial x_k} g_k(x_1^{n+1}, \dots, x_{k-1}^{n+1}, x_k^n, \dots, x_M^n)} \quad (4-27)$$

The term, ω , is a scalar typically set between $0 < \omega < 2$. Normally, each iteration of the NLOR technique requires evaluation of the entire objective function and its derivative. However, for the PET deformation problem, a small change in a single motion vector component affects only a very small fraction of all the terms in the function. Therefore the NLOR technique can be carried out relatively efficiently. In our implementation, we initialize the motion field to zero, then use a checkerboard update to proceed with the iteration. Additionally, we found that convergence was improved if a limit equal to the voxel size was imposed on the maximum step size.

Chapter 5

Deformation Sampling

5.1. Forward Verses Backward Sampling

Though the motion field describing the volume deformation is a one-to-one, onto mapping in a continuous domain, implementation on a discrete domain involves some subtleties that are important to recognize in the deformation of PET data sets. In an application where the deformed volumes are being summed together to form a composite volume, care must be taken during deformation sampling so that the composite volume represents a sum of all voxels from the source volumes. Past efforts implementing 3D deformations such as the work of Christensen, Bajcsy, and Zhou [15,7,103] have used a reverse transformation to deform voxels in the deformed volume. In this Eulerian formulation, the flow field motion vectors describe a particle’s motion with respect to its final position. That is, the motion vectors are defined in the space of the reference volume, and the “heads” of all the motion vectors terminate at discrete voxel locations in the reference volume. The motion vector “tails” effectively sample from a continuously defined location in the source volume. Thus to obtain the value of each voxel in the deformed volume, $\hat{f}(\mathbf{r}) = f_1(\mathbf{r} - \mathbf{m})$, eight voxels from the deformation volume are sampled at the location, $\mathbf{r} - \mathbf{m}$, and weighted according to trilinear interpolation. In the warping of two-dimensional scene, a backward sampling scheme may be useful for overcoming problem areas due to occlusion, where an occluded region in the source image is unoccluded in the reference volume. However, for three-dimensions continuous volumes, no occlusion problems exist, and each

voxel in a source volume should be represented in the reference. *Backward* sampling unfortunately does not guarantee that each voxel in the deformation volume will contribute to the deformed volume.

In a previous work [47], we proposed the use of a forward sampling technique which defines the location of the motion vector at its starting position in the space of the source volume. In this Lagrangian formulation, the motion of each discrete voxel in the source volume is well-defined, though the motion vector generally terminates on non-integer spatial locations within the deformed volume. To calculate the deformation, the value of each voxel in the source volume is simply distributed into the eight voxels at the location, $\mathbf{r} + \mathbf{m}$, via trilinear interpolation. The deformed volume is calculated in a single pass by first initializing the volume to zero, then forward projecting all voxels in the source volume, adding to the appropriate eight voxels in the deformed volume for each projection. Formally, this may be expressed by

$$\hat{f}(\mathbf{r}) = \sum_{\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}) \in R(\mathbf{r})} w_{\tilde{\mathbf{r}}} f_1(\tilde{\mathbf{r}}) \quad (5-1)$$

where $w_{\tilde{\mathbf{r}}}$ represents the trilinear weighting factor and $R(\mathbf{r})$ represents the region into which a forward displaced voxel would contribute.




Figure 5-1. Deformation Sampling Techniques

Forward sampling uses Lagrangian motion vectors to displace each voxel from a source volume into target voxels of the deformed volume. In contrast, backward samples use motion vectors defined in an Eulerian sense. Depending on the motion field, some voxels may contribute little or not at all to the deformed volume. The forward sampling technique insures that each voxel in the source volume is represented in the deformed volume.

This forward deformation technique guarantees that each voxel in the target volume is represented in the deformed volume and that the integral of voxel intensities is conserved between the source and deformation volumes. Additionally, it has the characteristic that a converging flow field sampling from a uniform region in the source volume results in increased voxel intensities for corresponding regions in the deformed volume. In a similar manner, a diverging flow field produces corresponding voxel intensity decreases. Though it would seem that this would be a useful characteristic when modeling incompressible underlying media, it was found that in practice artifacts frequently developed in the deformed volume during the calculation of an optimal flow field. This could be attributed to two things. First, even small non-zero divergence in the motion field produces nonuniformities in the deformed volume due to a nonuniform redistribution of source voxel intensities into the target deformation volume. As such, it was difficult for a deformable motion minimization algorithm to converge to a motion field producing a deformation matching the reference volume. Second, the forward deformation calculated in this fashion is a non-linear function of the motion vectors. Hence, the validity of a Taylor series approximation of $\hat{f}(\mathbf{r})$ as described previously ($\hat{f}(\mathbf{r}) = f_1(\mathbf{r} + \tilde{\mathbf{m}}) - \nabla f_1(\mathbf{r} + \tilde{\mathbf{m}})\delta\mathbf{m}$) is questionable.

A modified implementation used in this work makes use of a normalized Gaussian weighted sampling scheme. Like the forward sampling scheme described by Equation 5-1, this scheme employs a motion field defined in a Lagrangian sense. Recognizing that the motion vectors will terminate at non-discrete locations in the deformation space, the contribution that each source voxel makes to a voxel in the deformation volume is weighted by a Gaussian distance measure. The distance measure is subsequently normalized by the total weights of all source voxels contributing to each target voxel. More exactly, the value of the deformed volume at some discrete location on the voxel grid, \mathbf{r} , may be expressed as

$$\hat{f}(\mathbf{r}) = \frac{\sum_{d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r}) \leq D_{max}} w_{d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r})} f_1(\tilde{\mathbf{r}})}{\sum_{d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r}) \leq D_{max}} w_{d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r})}} \quad (5-2)$$

where $d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r})$ represents the Euclidean distance between the target voxel in the deformation volume at location \mathbf{r} and the location of the displaced source voxel at

$\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}})$, and where D_{max} represents the maximum distance from a target voxel into which a displaced source voxel can contribute. The weighting factor, $w_{d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r})}$, is given by the Gaussian function as follows:

$$w_{d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r})} = \frac{1}{\sqrt{2\sigma}} e^{-\frac{(d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r}))^2}{2\sigma^2}} \quad (5-3)$$

This scheme is similar to a technique called kriging [68], which is often used to obtain samples on a regular grid from a set of non-uniformly spaced samples. Like our previous formulation, the forward deformation is implemented in a single pass by first initializing the deformed volume to zero, then forward projecting each voxel in the target volume. At the location of the forward projection, the distance is calculated to neighboring voxels in the resampled volume and the projected voxel is distributed according to the calculated weighting. By storing the total weights accumulated for each resampled voxel, a normalization factor can be calculated at the end of the pass that insures that all weights at each target voxel add to one.




Figure 5-2. Ellipsoidal Phantom.

Three orthogonal views (transverse, coronal, sagittal) through a 3D ellipsoidal phantom. For reference, each image shows the lines of intersection with the two other views.

5.2. Sampling Examples

A hollow ellipsoid phantom, as seen in Figure 5-2 will be used to illustrate some the characteristics of the different sampling schemes. Here, the ellipsoid is represented in a $64 \times 64 \times 16$ voxel density volume as a bright object on a zero background. The figure shows three orthogonal views to emphasize that it is a 3D phantom, though only the transverse view will be shown in subsequent figures. To first demonstrate that in some cases, nearly




Figure 5-3. Simple Translation.

Uniform translation using different sampling schemes produces indistinguishable results. Slice through the source volume is seen in (a), the vector flow field in (b), and the “true” translated deformed image in (c). All three sampling schemes produce indistinguishable results, shown in (d). However, due to the fact that the backward scheme does not necessarily sample from all source voxels, a great number of “incorrect” flow fields, as shown in (e) may also produce a similar deformation (f). Using the same “incorrect” flow field in a forward scheme produces a predictable “incorrect” deformed volume, as in (g).

any naive sampling scheme is adequate, Figure 5-3 shows the results after deforming the volumes via a uniform translation motion field. A transverse slice through the source volume is seen in (a), and the uniform translation flow field along this slice is seen as a vector plot in (b). The “truly” translated ellipsoid is seen in (c) along with an edge map of the original source volume. The results of the deformation using a backward sampling scheme, an unnormalized forward sampling scheme, and a Gaussian normalized scheme all produce the same result, as seen representatively in (d). A weakness of the backward sampling technique is demonstrated by using the extremely non-uniform flow field seen in (e). This vector field represents an extremely divergent case where nearly all motion vectors point back to just a few source voxels. The resulting deformation, seen in (d), is exactly the same as for the translational case. Conversely, when this motion field is applied to




Figure 5-4. Uniform Rotation.

Motion field describing a rigid rotation about the z-axis is seen in (a). Resulting deformation volumes obtained using backward sampling (b) and normalized forward sampling (c) appear accurate, but the volume obtained using unnormalized forward sampling has a noisy appearance due to the nonuniform redistribution of source voxels. Grey scales in (b),(c), and (d) have all been normalized by the overall minimum and maximum voxels in found in each volume.

either of the forward sampling schemes, the result is obviously different from a uniform translation, as seen in (g). This example points out how backward sampling can under-represent a considerable number of voxels from the source volume, and how the domain of permissible motion fields which produce a deformation matching some reference volume can be quite large. When using a forward sampling scheme, the range of motion fields that could produce acceptable deformed images matching a reference would be more limited.

The next example illustrates where a normalization term is useful in a forward sampling scheme. Figure 5-4 shows the results of the ellipsoid undergoing a uniform 30 degree rotation. Motion field vectors are seen in (a), and the resulting deformation using backward and normalized forward sampling are seen in (b) and (c) respectively. These two images show nearly identical results, that is, a rotated version of the source ellipse that is only different because of the intrinsic smoothing of the interpolation process. In contrast, the result using unnormalized forward sampling, seen in (d) looks markedly different. Here, the discrete implementation of deformation sampling with a continuous rotational motion field results in a mottled deformed image due to a slightly uneven redistribution of source voxels. In initial efforts to implement deformable motion algorithms using this technique, it was found that convergence of the algorithm was extremely slow. Presumably, this could be due to the algorithm attempting to minimize image differences due solely to this non-

uniform redistribution, rather than due to actual motion between a source and reference volume.

A final example using an incompressible stretching deformation is seen in Figure 5-5. The images in (a) and (b) show the source volume before and after stretching the volume by factors of 1.25 and 0.8 respectively along the x and y axes respectively. Two different interpretations of the motion field implementing this deformation are shown in (c) and (d). The first of these is a uniform extension and compression in X and Y along the entire volume. The second is a discontinuous motion field that is non-zero only at locations corresponding to the voxels representing the source ellipsoid. Note that in this case, the discontinuous motion field could not be used to describe the motion of a continuous volume, such as a cube of solid rubber, for it violates the principle of the impenetrability of matter. However, it may make sense to consider a motion description like this in some cases, for example for a shearing motion between two adjacent but not connected surfaces, or for the motion of an object next to a featureless background, as in the case of the brightly defined left ventricle in cardiac PET data. The images seen in (e), (f) and (g) show the results of using the backward, normalized forward and unnormalized forward sampling schemes on the continuous motion field in (c). These three images show similar results to the uniform rotation case. That is, both backward and normalized forward sampling produce a deformed volume representative of the true deformation. Again, the volume produced by unnormalized forward sampling matches the shape of the desired deformed ellipsoid, but the voxel values are nonuniform due to the nonuniform redistribution of the source voxels.

The last row of images shows the results of the three sampling schemes on the discontinuous motion field. The inadequacies of the backward sampling case are obvious, since the shape of the deformed ellipsoid is severely distorted in (h). The normalized forward sampling technique produced a better estimate, preserving the shape of the stretched ellipsoid even though the voxel intensities at the edges are somewhat in error due to the smoothing inherent to the Gaussian weight of displaced voxels. Only the unnormalized forward sampling technique conserves the integral of voxel intensities between the source and the deformed images, and appears to preserve the shape transformation as well. However,




Figure 5-5. Incompressible Stretch Deformation.

Source volume (a) and stretched reference (b). A continuous motion field implement the stretch is seen in (c); the same field with positive motion vectors only at locations corresponding to the ellipsoid (d). Deformations calculated using backward (e), normalized forward (f) and unnormalized forward (g) sampling and the continuous motion field of (c). The correspond deformations using the discontinuous motion field of (d) for the three sampling techniques are shown in (h), (i) and (j).

again in this case the resulting deformation in (j) is nonuniform in areas where the object should be uniform. For these last three data sets, the ratio of the actual integral of voxel intensities to true integral of voxel intensities was 1.15, 0.96 and 1.0 for (h), (i) and (j); and the squared error difference between the reference in (b) and the three are 8.68148×10^7 , 4.97173×10^7 , and 1.75321×10^7 . This compares to voxel intensity ratios of 0.999, 1.0,

0.999 and squared error differences of $1.34487\text{e}+07$, $1.45759\text{e}+07$, $1.7442\text{e}+06$ for the volumes (e), (f) and (g) produced by the continuous motion field. So it would appear that for continuously defined incompressible motion fields, either a backward or a forward sampling technique are adequate, though a forward sampling technique may be more useful in disallowing physically implausible fields. For discontinuous motion fields, a forward sampling technique appears to be the most faithful method.

Chapter 6

Anisotropic Elastic Model

6.1. Stress and Strain Relations

In order to develop a model for material strain energy, it is necessary to review the basic relations between stress and strain in a setting for both infinitesimal displacements and finite displacements. Unless otherwise cited, all formulations here have been derived from the classical books on elasticity by Love and Washizu [54, 96]. In the simplest of terms, stress is defined as a distributed force that is applied to a solid body, and strain describes the resulting relative displacement of the body in response to

the stress. There are nine components to the stress tensor defining the state of internal force at a point in a body. These include three longitudinal stresses: σ_{xx} , σ_{yy} , σ_{zz} , and six tangential stresses: $\sigma_{xy} = \sigma_{yx}$, $\sigma_{zy} = \sigma_{yz}$, $\sigma_{xz} = \sigma_{zx}$. Shown in Figure 6-1, for stresses acting on an infinitesimal rectangular parallelepiped, the longitudinal stresses are forces




Figure 6-1. Stress Tensor

Six independent components of the strain tensor acting on a infinitesimal rectangular parallelepiped

that induce a stretch along each of the principle axes, and the tangential stresses induce a shearing. As these forces deform an elastic body, we can define six independent components of strain to describe the relative displacement: $e_{xx}, e_{yy}, e_{zz}, e_{xy}, e_{yz}, e_{xz}$. Like the components of the stress tensor, the extensional strain components, e_{xx}, e_{yy}, e_{zz} describe the change in length, or the stretching, for any linear element oriented along the principle axes; and the shear strain components, e_{xy}, e_{yz}, e_{xz} , describe the change in angles between any two linear elements. In small displacement theory, the relation between the stress and strain is given by this simple linear form, known as the generalized Hooke's Law:

$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{zz} \\ \sigma_{yz} \\ \sigma_{xz} \\ \sigma_{xy} \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} \\ a_{31} & a_{32} & a_{33} & a_{34} & a_{35} & a_{36} \\ a_{41} & a_{42} & a_{43} & a_{44} & a_{45} & a_{46} \\ a_{51} & a_{52} & a_{53} & a_{54} & a_{55} & a_{56} \\ a_{61} & a_{62} & a_{63} & a_{64} & a_{65} & a_{66} \end{bmatrix} \begin{bmatrix} e_{xx} \\ e_{yy} \\ e_{zz} \\ e_{yz} \\ e_{xz} \\ e_{xy} \end{bmatrix} \quad (6-1)$$

For an isotropic material, only two independent elastic constants exist, and the stress-strain relation reduces to:

$$\begin{aligned} \sigma_{xx} &= 2\mu \left[e_{xx} + \frac{\nu}{1-2\nu} (e_{xx} + e_{yy} + e_{zz}) \right] \\ \sigma_{yy} &= 2\mu \left[e_{yy} + \frac{\nu}{1-2\nu} (e_{xx} + e_{yy} + e_{zz}) \right] \\ \sigma_{zz} &= 2\mu \left[e_{zz} + \frac{\nu}{1-2\nu} (e_{xx} + e_{yy} + e_{zz}) \right] \end{aligned} \quad \begin{aligned} \sigma_{yz} &= \mu e_{yz} \\ \sigma_{xz} &= \mu e_{xz} \\ \sigma_{xy} &= \mu e_{xy} \end{aligned} \quad (6-2)$$

where ν is Poisson's ratio and μ is the Lamé constant defined previously in Chapter 4, also known as the modulus of rigidity. (Recall that $\nu = \lambda / (2(\lambda + \mu))$.)

A more intuitive feel for the meaning of the strain components as well as the distinction between the finite displacement and small displacement (i.e. infinitesimal) strain components may be obtained by observing the parallelepiped volume element in Figure 6-2




Figure 6-2. Volume Element Undergoing Deformation

Material points in an undeformed volume element are deformed into a new configuration.

before and after deformation. Coordinates for points in the undeformed volume, for example, $\mathbf{R} = (R_1, R_2, R_3)$, are said to be in a *material* coordinate system aligned to the principle axes of a rectangular Cartesian coordinate system. Each point in the undeformed volume is displaced by a vector, $\mathbf{m} = (u, v, w)$ to a new position in the deformed volume, $\mathbf{r} = (r_1, r_2, r_3)$ with coordinates specified in the material coordinate system. (We may also define another coordinate system, called the *reference* coordinate system that deforms with the body.) Thus, we have the relation, $\mathbf{r} = \mathbf{R} + \mathbf{m}$. The deformation gradient tensor is defined in terms of the material coordinate system as follows:

$$\mathbf{F} = \left[\frac{\partial r_i}{\partial R_j} \right] \quad (6-3)$$

Next, the displacement gradient tensor may be defined in terms of the displacement vector, \mathbf{m} , as the gradient of the motion field,

$$\nabla \mathbf{m} = \mathbf{U} = \begin{bmatrix} u_x & u_y & u_z \\ v_x & v_y & v_z \\ w_x & w_y & w_z \end{bmatrix} \quad (6-4)$$

Note that the displacement field, \mathbf{m} , is synonymous with our Lagrangian motion field from Chapter 4. This gradient tensor may be used to define the right Cauchy-Green tensor, also called the Green deformation tensor,

$$\mathbf{C} = \mathbf{F}^T \mathbf{F} = \mathbf{I} + \mathbf{U} + \mathbf{U}^T + \mathbf{U}^T \mathbf{U} \quad (6-5)$$

Finally, the Green-Lagrange (or Lagrangian) *strain* tensor is defined by

$$\mathbf{E} = \frac{1}{2}(\mathbf{C} - \mathbf{I}) = \frac{1}{2}(\mathbf{U} + \mathbf{U}^T + \mathbf{U}^T \mathbf{U}) \quad (6-6)$$

We may now use Figure 6-2 and a few examples to lend some intuition to the physical interpretation of these tensors. It is immediately apparent from Equation 6-3 that \mathbf{F} describes the deformation of a small line, $d\mathbf{x} = \mathbf{F}d\mathbf{X}$. Consider now a simple dilation which stretches the parallelepiped by the factors, $\alpha_1, \alpha_2, \alpha_3$, along the principle axes. The deformation gradient tensor and Cauchy-Green tensors corresponding to this deformation are

$$\mathbf{F} = \begin{bmatrix} \alpha_1 & 0 & 0 \\ 0 & \alpha_2 & 0 \\ 0 & 0 & \alpha_3 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} (\alpha_1)^2 & 0 & 0 \\ 0 & (\alpha_2)^2 & 0 \\ 0 & 0 & (\alpha_3)^2 \end{bmatrix} \quad (6-7)$$

So for the simple case of pure dilation, the diagonal components of the Cauchy-Green tensor relate the square of the elongation along each axis. That is, for example, for an undeformed line segment, $d\mathbf{R}$ aligned to the X_2 axis as shown in the figure, the rate of elongation between $d\mathbf{R}$ and the corresponding line segment, $d\mathbf{r}$, is given by $(d\mathbf{r} - d\mathbf{R})/d\mathbf{R} = \sqrt{C_{22}} - 1$. In fact, this trait is true in general for the three diagonal com-

ponents of \mathbf{C} even when deformed line segments are no longer aligned to the principle axes. This is demonstrated in the following case of simple shear. A shearing along the X_3 axis by an angle, $\theta = \text{atan} \kappa$ is described by the following deformation gradient and Cauchy-Green tensors

$$\mathbf{F} = \begin{bmatrix} 1 & \kappa & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 1 & \kappa & 0 \\ \kappa & 1 + \kappa^2 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (6-8)$$

Figure 6-3 shows the deformation for this example. It is seen that though $d\mathbf{r}$ is no longer aligned to the X_2 axis, the elongation relation holds. The example also points out the geometrical meanings of the off-diagonal elements of the




Figure 6-3. Simple Shear Example

Cauchy-Green tensor. These components describe the change in angle between two line segments in the undeformed and deformed volumes. For example, the change in angles between $d\mathbf{R}$ (aligned to the X_2 axis) and $d\tilde{\mathbf{R}}$ aligned to the X_1 axis compared with $d\mathbf{r}$ and $d\tilde{\mathbf{r}}$ is given by

$$C_{12} = \sqrt{C_{11}C_{22}} \sin \theta_{12} \quad (6-9)$$

where in general $(\pi/2) - \theta_{12}$ is the angle between the two deformed line segments originally oriented along the X_1 and X_2 axes respectively. Similar relations follow in general for C_{23} and C_{13} .

Now let us consider the Green-Lagrange strain tensor. If we expand the components of this tensor using Equation 6-6 and the components of the displacement gradient tensor, we have

$$\mathbf{E} \equiv \begin{bmatrix} \epsilon_{11} & \epsilon_{21} & \epsilon_{31} \\ \epsilon_{12} & \epsilon_{22} & \epsilon_{23} \\ \epsilon_{13} & \epsilon_{23} & \epsilon_{33} \end{bmatrix} \quad (6-10)$$

where

$$\begin{aligned} \epsilon_{11} &= u_x + \frac{1}{2}[u_x^2 + v_x^2 + w_x^2] \\ \epsilon_{22} &= v_y + \frac{1}{2}[u_y^2 + v_y^2 + w_y^2] \\ \epsilon_{33} &= w_z + \frac{1}{2}[u_z^2 + v_z^2 + w_z^2] \\ \epsilon_{12} = \epsilon_{21} &= \frac{1}{2}[u_y + v_x + (u_x u_y + v_x v_y + w_x w_y)] \\ \epsilon_{23} = \epsilon_{32} &= \frac{1}{2}[v_z + w_y + (u_y u_z + v_y v_z + w_y w_z)] \\ \epsilon_{13} = \epsilon_{31} &= \frac{1}{2}[u_z + w_x + (u_x u_z + v_x v_z + w_x w_z)] \end{aligned} \quad (6-11)$$

This is what is known as the finite displacement strain tensor, for it is a valid description of material strain for both small and large displacements. We distinguish this from the infinitesimal (small displacement) strain tensor:

$$\mathbf{e} \equiv \begin{bmatrix} e_{11} & e_{21} & e_{31} \\ e_{12} & e_{22} & e_{23} \\ e_{13} & e_{23} & e_{33} \end{bmatrix} \quad (6-12)$$

where the second degree terms are assumed to be negligible and are omitted:

$$\begin{aligned} e_{11} &= u_x, e_{22} = v_y, e_{33} = w_z \\ e_{12} = e_{21} &= \frac{1}{2}[u_y + v_x] \\ e_{23} = e_{32} &= \frac{1}{2}[v_z + w_y] \\ e_{13} = e_{31} &= \frac{1}{2}[u_z + w_x] \end{aligned} \quad (6-13)$$

The geometrical interpretation of the large displacement strain tensors parallels those of the Cauchy-Green tensor. The diagonal elements of the tensor describe object stretching along the principle axes, and the off-diagonal elements describe changes in angles between two line segments. More exactly, these relations hold [44]:

$$\begin{aligned} (d\mathbf{r}_1 - d\mathbf{R}_1)/d\mathbf{R}_1 &= \sqrt{1 + 2\varepsilon_{11}} - 1 \approx \varepsilon_{11} \\ \theta_{12} &= \text{asin}\left(\frac{2\varepsilon_{12}}{\sqrt{1 + 2\varepsilon_{11}}\sqrt{1 + 2\varepsilon_{22}}}\right) \approx 2\varepsilon_{12} \end{aligned} \quad (6-14)$$

for a line segment $d\mathbf{R}_1$ aligned to the X_1 axis in the undeformed volume, and for a change in angle, θ_{12} between two line segments initially aligned along the X_1 and X_2 axes. The approximation on the right is valid for small displacements. The relations for the other strain components follow analogously.

6.2. Strain Energy Functions

As an elastic body is deformed, the work that is done by the external forces results in stored potential energy for that body. Assuming that the deformation does not change the temperature of the object, this stored energy is called the *strain energy* associated with the deformed body. Washizu shows that using small displacement theory, the strain energy stored in a unit volume is given by

$$A(e_{11}, \dots, e_{13}) = \int_{(0,0,0,0,0,0)}^{(e_{11}, e_{22}, e_{33}, e_{23}, e_{13}, e_{12})} (\sigma_{11}de_{11} + \dots + \sigma_{12}de_{12}) \quad (6-15)$$

Assuming that the stress-strain relations are linear, as in Equation 6-1, and $\sigma_{ij}de_{ij}$ is a perfect differential that satisfies the relation,

$$\frac{\partial \sigma_{\lambda\mu}}{\partial e_{\alpha\beta}} = \frac{\partial \sigma_{\alpha\beta}}{\partial e_{\lambda\mu}}, (\lambda, \mu, \alpha, \beta = 1, 2, 3) \quad (6-16)$$

then the strain energy equation can be simplified to

$$A = \frac{1}{2} \{\mathbf{e}\}^T \mathbf{a} \{\mathbf{e}\} \quad (6-17)$$

where $\{e\}^T$ is the vector of infinitesimal strain components, $\{e_{11}, e_{22}, e_{33}, e_{23}, e_{13}, e_{12}\}$ and \mathbf{a} is the matrix relating stress and strain in Equation 6-1. A similar expression may be found using finite displacement theory, except that the infinitesimal strain components are replaced with the finite theory versions, $\{\epsilon_{11}, \epsilon_{22}, \epsilon_{33}, \epsilon_{23}, \epsilon_{13}, \epsilon_{12}\}$. For isotropic materials, the strain energy function simplifies to

$$A = \frac{\lambda}{2}(e_{11} + e_{22} + e_{33})^2 + \mu(e_{11}^2 + e_{22}^2 + e_{33}^2) + 2\mu(e_{12}^2 + e_{13}^2 + e_{23}^2) \quad (6-18)$$

and for the finite displacement model, the strain energy function is

$$A = \frac{\lambda}{2}(\epsilon_{11} + \epsilon_{22} + \epsilon_{33})^2 + \mu(\epsilon_{11}^2 + \epsilon_{22}^2 + \epsilon_{33}^2) + 2\mu(\epsilon_{12}^2 + \epsilon_{13}^2 + \epsilon_{23}^2) \quad (6-19)$$

Note that the strain energy for the infinitesimal model is equivalent to the linear elastic constraint already described by Equation 4-14.

Numerous other strain energy relations have been proposed to describe linearly deforming and non-linearly deforming materials. Many linear models make use of the three principle invariants of the Cauchy-Green tensor, \mathbf{C} , which remain unchanged under coordinate rotations. These are given by [5,43]:

$$\begin{aligned} I_1 &= \text{tr}\mathbf{C} = (\alpha_1)^2 + (\alpha_2)^2 + (\alpha_3)^2 \\ I_2 &= \frac{1}{2}[(\text{tr}\mathbf{C})^2 - \text{tr}\mathbf{C}^2] = (\alpha_1)^2(\alpha_2)^2 + (\alpha_2)^2(\alpha_3)^2 + (\alpha_1)^2(\alpha_3)^2 \\ I_3 &= \det\mathbf{C} = (\alpha_1)^2(\alpha_2)^2(\alpha_3)^2 \end{aligned} \quad (6-20)$$

where $\text{tr}\mathbf{C}$ denotes the trace of the \mathbf{C} matrix, $\det\mathbf{C}$ denotes the determinant of \mathbf{C} , and where $\alpha_1, \alpha_2, \alpha_3$ are the components of stretch along the three principle axes. Because a description of an isotropic material must be invariant to a change of coordinates, all isotropic linear strain energy models may be expressed in terms of these invariants. If the material is also incompressible, then $\alpha_1\alpha_2\alpha_3 = 1$, so $I_3 = 1$, and the strain energy functions for these materials can be expressed in terms of just the first two invariants. One strain energy function, called the “neo-Hookean” model is given by:

$$A = C_1(I_1 - 3) = 2C_1(\epsilon_{11} + \epsilon_{22} + \epsilon_{33}) \quad (6-21)$$

This energy model has been found to characterize materials like vulcanized rubber quite well. Another related model which is also used to model incompressible, rubber-like materials is the so-called Mooney-Rivlin form:

$$A = C_1(I_1 - 3) + C_2(I_2 - 3) = 2C_1(\epsilon_{11} + \epsilon_{22} + \epsilon_{33}) + 4C_2[(\epsilon_{11} + \epsilon_{22} + \epsilon_{33}) + (\epsilon_{11}\epsilon_{22} + \epsilon_{22}\epsilon_{33} + \epsilon_{11}\epsilon_{33})] \quad (6-22)$$

When the constant weighting parameters, C_1 and C_2 were chosen properly, this second model was seen to more accurately predict experimental data from certain kinds of rubber than did the neo-Hookean form [5]. At first glance, these functions look considerably different from the linearly elastic model of Equation 6-18. However, Rivlin has shown that at least the neo-Hookean form is equivalent to the linearly elastic model for incompressible materials under small strains [76]. Because of the limitation to completely incompressible materials and the small strain requirement, these simpler models are probably not appropriate for modeling cardiac tissue, and they are certainly not as general as the form in either Equation 6-18 or Equation 6-19.

Accurately modeling the strain energy function for biological tissue is a difficult task. It is known that tissue such as the myocardium is anisotropic. Elastic properties vary in the different directions due to muscle fiber orientations and other factors. Also soft biological tissue often responds fairly nonlinearly and with large deformations. In addition, though the tissue is nearly incompressible, there is at least some change in volume due to different volumes of blood being present in the capillary bed and other small vessels embedded in the tissue. Some authors have chosen to accurately model this phenomenon, incorporating information about typical muscle fiber orientations, anisotropic tissue elasticity properties, and other regional effects that require alignment of a heart data set with a well-described cardiac template [66,43]. Once this is done, a strain energy such as the following is possible:

$$A = k_1 \frac{e_{11}^2}{(a_1 - |e_{11}|)^{\beta_1}} + k_2 \frac{e_{22}^2}{(a_2 - |e_{22}|)^{\beta_2}} + k_3 \frac{e_{33}^2}{(a_3 - |e_{33}|)^{\beta_3}} + k_4 \frac{e_{12}^2}{(a_4 - |e_{12}|)^{\beta_4}} + k_5 \frac{e_{23}^2}{(a_5 - |e_{23}|)^{\beta_5}} + k_6 \frac{e_{13}^2}{(a_6 - |e_{13}|)^{\beta_6}} \quad (6-23)$$

where e_{ij} are the components of the Green strain tensor in the coordinate system measured with respect to the tissue, and the other factors are constants used to express the anisotropic nonlinear behavior along each axis. Strain energy models like these can be quite effective in characterizing cardiac tissues, especially when combined with finite element models of the heart structure. However, this type of model is extremely expensive computationally, and requires an accurate segmentation of the cardiac tissue in a 3D density volume, as well as recognition of landmarks to establish proper registration with cardiac models [20]. It is hoped that for the deformable motion problem here, adequate results can be obtained with simpler models.

6.3. Cardiac PET Piece-wise Isotropic Strain Energy Model

Up until this point, we have outlined various models with a wide range of complexity and accuracy for characterizing biological tissue in order to regularize the deformation problem in cardiac PET. However, we have not yet addressed the fact that dramatically different tissue types are present within the field of view. Indeed, the elastic properties of cardiac tissue are anisotropic, as modeled in Equation 6-23, but the anisotropic properties of cardiac tissue compared with the blood pool inside the ventricles, or with the adjacent lung tissue and air space are bound to be a larger effect. It is for this reason that we desire to model this latter type of anisotropic elasticity, while still utilizing an algorithm that does not require complex pattern recognition of specific cardiac features.

We propose a piece-wise isotropic strain energy model which models the cardiac tissue separately from the blood pool and other tissue types in the reconstructed PET volume. By using an isotropic model for each tissue type, the amount of pre-processing, such as segmentation and landmark recognition, is kept to a minimum. Furthermore, a linear isotropic strain energy model is fairly easy to implement in a minimization scheme, and its corresponding computational cost does not become burdensome. For the bulk of the remaining

work, we will use a modification of the linear strain model described by Equation 6-18 as follows:

$$A(\mathbf{r}) = \frac{\lambda(\mathbf{r})}{2}(e_{11}(\mathbf{r}) + e_{22}(\mathbf{r}) + e_{33}(\mathbf{r}))^2 + \mu(\mathbf{r})(e_{11}^2(\mathbf{r}) + e_{22}^2(\mathbf{r}) + e_{33}^2(\mathbf{r})) + 2\mu(\mathbf{r})(e_{12}^2(\mathbf{r}) + e_{13}^2(\mathbf{r}) + e_{23}^2(\mathbf{r})) \quad (6-24)$$

Thus, the linear elastic model has been generalized to allow varying elastic weighting parameters over the volume, $\lambda(\mathbf{r})$ and $\mu(\mathbf{r})$ instead of just two global scalars. For locations corresponding to the nearly incompressible tissue of the heart, the Lamé constants may be set such that the Poisson ratio approaches 0.5. Conversely, for locations such as the blood which should not be modeled as incompressible (or possibly, even conserved), the $\lambda(\mathbf{r})$ term may be made small, and only the $\mu(\mathbf{r})$ term would be used to promote a smooth motion field in those areas.

In this dissertation, we will only consider a bi-valued vector field giving the variable Lamé constants: one set of values for voxels labeled as cardiac tissue, and another set of values for voxels labeled as background. The preprocessing requirement thus requires a segmentation of the left ventricle from the background. We note that an accurate automatic segmentation of the myocardial boundaries in a gated cardiac PET volume may be a formidable task itself, especially due to the noisy nature of the data and the partial volume effects due to spatial resolution limits. For the results in this dissertation, manual segmentations will be used. Because this segmentation is only being used to model differing elastic characteristics of different tissue types, and not to define different sampling characteristics, it should not be sensitive to slight errors of a few voxels along the cardiac boundaries. This is in contrast to a technique that would pre-segment voxels corresponding to cardiac tissue, and only operate on (i.e. displace via the deformation algorithm) those voxels.

We acknowledge that the infinitesimal description assumes that the displacement components are small, and thus this description is an approximation for typical motion of the heart. Waldman [95] has shown that because of the large-distance deformations of the heart during the cardiac cycle, the infinitesimal approximation can induce errors of at least 16% in the strain components. However, because a large-distance strain energy formula-

tion involves second degree terms of the motion field partial derivatives, an algorithm implementing a minimization procedure is more computationally intensive and perhaps is more sensitive to noise in the partial derivative estimates. For this reason, we choose to use the infinitesimal approximation for most of the results in this dissertation, and note that its smoothing characteristics should be similar even though the actual strain energies calculated may be quite different. A comparison of the finite displacement and infinitesimal models is found in the results of this work.

Chapter 7

Final Formulation

7.1. Algorithm Summary

To summarize the different facets of the deformable motion technique described thus far, the goal of the deformable motion algorithm is to calculate a dense motion vector field so that a correspondence may be made between matching voxels in a source and a reference volume data set. It is assumed that the source and reference data set voxels represent the same physical object that has been non-rigidly deformed between volume imaging times. The motion field is used along with the source volume to compute a deformed volume with features that closely match features in the reference volume. Deformation of the source volume is implemented using a normalized forward sampling technique defined by Equation 5-2, and repeated here for convenience:

$$\hat{f}(\mathbf{r}) = \frac{\sum_{d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r}) \leq D_{max}} w_{d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r})} f_1(\tilde{\mathbf{r}})}{\sum_{d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r}) \leq D_{max}} w_{d(\tilde{\mathbf{r}} + \mathbf{m}(\tilde{\mathbf{r}}), \mathbf{r})}} \quad (7-1)$$

While computing the best motion field to deform the source volume, two criteria are weighed against one another; first an image matching criterion which tries to minimize a difference between the deformed and the reference volumes; and second, a motion field smoothness criterion, which prevents the motion field from taking configurations that would be impossible or unlikely for an actual physical deforming object. The squared difference between single voxels is used as a measure of goodness for the image matching

criterion, and a piece-wise isotropic linear elastic strain energy function is used a measure of goodness for the regularity criterion. Therefore, the motion field is obtained by minimizing the following cost function:

$$E_{tot} = \sum_{\mathbf{r}} [e_I(\mathbf{r}) + e_S(\mathbf{r})] \quad (7-2)$$

where $e_I(\mathbf{r})$ is the image matching measure given by Equation 4-7, shown again here:

$$e_I(\mathbf{r}) = \gamma_I(f_2(\mathbf{r}) - \hat{f}(\mathbf{r}))^2 \quad (7-3)$$

and $e_S(\mathbf{r})$ is the piece-wise linear elastic strain energy given by either the small (infinitesimal) displacement strain assumption defined in Equation 6-24, and repeated here as:

$$\begin{aligned} e_S(\mathbf{r}) = & \frac{\lambda(\mathbf{r})}{2} [e_{11}(\mathbf{r}) + e_{22}(\mathbf{r}) + e_{33}(\mathbf{r})]^2 + \\ & \mu(\mathbf{r}) [e_{11}^2(\mathbf{r}) + e_{22}^2(\mathbf{r}) + e_{33}^2(\mathbf{r})] + 2\mu(\mathbf{r}) [e_{12}^2(\mathbf{r}) + e_{13}^2(\mathbf{r}) + e_{23}^2(\mathbf{r})] \end{aligned} \quad (7-4)$$

or by the finite displacement strain assumption, where the infinitesimal strain components, e_{ij} , in Equation 7-4 are replaced by the finite strain components, ϵ_{ij} , both which were defined in Chapter 6. Characteristics of different tissue types may be expressed in this strain energy equation by setting the elastic constants at each voxel location, $\lambda(\mathbf{r})$ and $\mu(\mathbf{r})$, to a value which is appropriate for that particular tissue.

When we choose the infinitesimal displacement strain formulation, the total cost function can be expressed in terms of the reference and deformed volumes along with the components of the motion field, $\mathbf{m}(\mathbf{r}) = (u(\mathbf{r}), v(\mathbf{r}), w(\mathbf{r}))$ as follows:

$$\begin{aligned} E_{tot} = & \sum_{\mathbf{r}} \gamma_I(f_2(\hat{\mathbf{r}}) - \hat{f}(\hat{\mathbf{r}}))^2 + \frac{\lambda(\mathbf{r})}{2} [u_x(\mathbf{r}) + v_y(\mathbf{r}) + w_z(\mathbf{r})]^2 + \\ & \mu(\mathbf{r}) [u_x^2(\mathbf{r}) + v_y^2(\mathbf{r}) + w_z^2(\mathbf{r})] + \\ & \frac{\mu(\mathbf{r})}{2} [u_y^2(\mathbf{r}) + u_z^2(\mathbf{r}) + v_x^2(\mathbf{r}) + v_z^2(\mathbf{r}) + w_x^2(\mathbf{r}) + w_y^2(\mathbf{r}) + \\ & 2u_y(\mathbf{r})v_x(\mathbf{r}) + 2u_z(\mathbf{r})w_x(\mathbf{r}) + 2v_z(\mathbf{r})w_y(\mathbf{r})] \end{aligned} \quad (7-5)$$

where the image matching term is calculated at the forward displaced location, $\hat{\mathbf{r}} = \mathbf{r} + \mathbf{m}(\mathbf{r})$. This equation is minimized by using a linear Taylor series approximation to the calculation of the deformed volume. As was described in Chapter 4, it is assumed that the current estimate of the motion field is $\tilde{\mathbf{m}}$, and is equal to the true motion field, \mathbf{m} , plus some small delta flow field, i.e. $\delta\mathbf{m} = \tilde{\mathbf{m}} - \mathbf{m}$. The calculus of variations is used to derive Euler-Lagrange equations and these are solved using a conjugate gradient descent method to find the best $\delta\mathbf{m}$.

It is interesting to note that for the infinitesimal strain case with non-varying Lamé constant values, the partial differential equations resulting from the calculus of variations (comparable to Equation 4-22, Equation 4-23, and Equation 4-24) can be rearranged in a format similar to Zhou's formulation [103], which expressed the regularization criteria in terms of a smoothness and a divergence-free constraint:

$$\gamma_I(f_2 - \hat{f} + \nabla \hat{f} \bullet \delta\mathbf{m})\hat{f}_x = (\lambda + \mu)(\delta u_{xx} + \delta v_{xy} + \delta w_{xz} - \tilde{u}_{xx} - \tilde{v}_{xy} - \tilde{w}_{xz}) + \mu(\delta u_{xx} + \delta u_{yy} + \delta u_{zz} - \tilde{u}_{xx} - \tilde{u}_{yy} - \tilde{u}_{zz}) \quad (7-6)$$

$$\gamma_I(f_2 - \hat{f} + \nabla \hat{f} \bullet \delta\mathbf{m})\hat{f}_y = (\lambda + \mu)(\delta u_{xy} + \delta v_{yy} + \delta w_{yz} - \tilde{u}_{xy} - \tilde{v}_{yy} - \tilde{w}_{yz}) + \mu(\delta v_{xx} + \delta v_{yy} + \delta v_{zz} - \tilde{v}_{xx} - \tilde{v}_{yy} - \tilde{v}_{zz}) \quad (7-7)$$

$$\gamma_I(f_2 - \hat{f} + \nabla \hat{f} \bullet \delta\mathbf{m})\hat{f}_z = (\lambda + \mu)(\delta u_{xz} + \delta v_{yz} + \delta w_{zz} - \tilde{u}_{xz} - \tilde{v}_{yz} - \tilde{w}_{zz}) + \mu(\delta w_{xx} + \delta w_{yy} + \delta w_{zz} - \tilde{w}_{xx} - \tilde{w}_{yy} - \tilde{w}_{zz}) \quad (7-8)$$

where we have again dropped the \mathbf{r} from the equations to simplify notation. That is, the Lamé stiffness coefficient, μ , is equivalent to Zhou's divergence-free weighting constant, and $(\lambda + \mu)$ is equivalent to Zhou's smoothness weighting constant. This points to the fact that to model physical materials, one cannot simply choose arbitrary values for the smoothness and divergence-free weights, since for physical materials, the Lamé constants are positive, and therefore $\mu \leq \lambda + \mu$.

When the finite displacement strain energy model is used as a regularization constraint, the cost function equation in terms of the motion field components comparable to Chapter 7-5 becomes quite unwieldy. Furthermore, it is no longer quadratic in $\delta\mathbf{m}$ if the

linear Taylor approximation is used to calculate the deformed volume. Therefore, for the minimization of this case, the nonlinear successive over-relaxation (NLOR) technique described in Chapter 4 is used.

7.2. Multi-Resolution Approach

Even though the image matching criterion used here is valid for large-scale displacements, convergence of either minimization technique to a desired minimal cost is difficult if the motion vectors are large with respect to the voxel size. Both minimization techniques make use of the gradient of the difference between the reference and the deformed volumes. The gradient is implemented using a standard central differences scheme, which is a local operator that is unaffected by voxel values more than one or two voxels away. Therefore, it would be difficult for the algorithm to begin “nudging” a motion vector from a value of zero if the displacement is much considerably larger than the voxel size.

A standard technique for dealing with problems of this nature is to use a scale space approach or a pyramid sampling scheme [98,51]. An approach like this employs a hierarchy of volumes obtained by subsampling the original volume at progressively coarser and coarser scales. Because the voxel size is larger at the coarser scales, larger displacements can be tracked using a minimizer that depends on image gradients. Typically, the coarser volumes are smoothed before subsampling to prevent aliasing while sparsely sampling the higher resolution volume. This avoids the generation of spurious detail upon increasing the scale. A common choice of the smoothing operator is a spherically symmetric Gaussian kernel. As the scale increases (and the sampling resolution decreases), the variance of the Gaussian kernel increases as well.

For the work presented here, we obtain a multi-resolution representation of the image volumes using a uniform cubic B-spline approximation to a Gaussian pyramid [75]. This is an efficient means of calculating a traditional scale space pyramid. As an example, typical cardiac PET volumes are $128 \times 128 \times 47$ voxel data sets with anisotropic voxel size of $2.0 \times 2.0 \times 3.1$ mm. This volume is subsampled into $64 \times 64 \times 23$, $32 \times 32 \times 23$ and $16 \times 16 \times 23$ data sets. The deformable motion algorithm is begun at the lowest resolution with the motion field initialized to zero. It is then run until convergence, and the resulting

motion field is used as an initial condition for the next higher scale. The initial motion field at a higher resolution is produced using trilinearly sampled interpolation of the lower resolution field.

The multi-resolution technique speeds the overall convergence, and in some cases where the displacements were quite large, the multi-resolution technique was found necessary to avoid solutions at incorrect local minima.

7.3. Calculating the Composite Volume

Once the deformed volume, $\hat{f}(\mathbf{r})$ matching $f_2(\mathbf{r})$ is obtained, subsequent processing to obtain a composite PET data set is straightforward. The composite sum is computed as

$$f_{sum}(\mathbf{r}) = \hat{f}(\mathbf{r}) + f_2(\mathbf{r}) \quad (7-9)$$

Because the deformed volume was producing using a forward mapping, each voxel from the source volume has been summed into the data set and the composite volume represents the total PET counts acquired in the two data sets. In general, data from all cardiac gates would be combined to form a single composite image, though the results presented here will only consider the summing of two frames.

Chapter 8

Results

8.1. Simulation Models

A problem with using real world medical imaging data to test an algorithm is that it is difficult to obtain a gold standard data set which establishes the “truth” volume. Past researchers have gone to the extreme of surgically implanting radioopaque markers in the epicardial surface of canine and human subjects in an attempt to accurately gauge cardiac motion in real cardiac data [94,59]. This usually is not possible, and for evaluating motion in PET data sets, a sparse array of markers would still be inadequate for quantifying the accuracy of a dense motion field calculation. For this reason, many of the results presented here will use simplified phantom data sets which model certain known characteristics of cardiac PET data. The first of these phantoms is a hollow ellipse that is deformed according to a uniform dilation using either a compressible or an incompressible material assumption. This will be used to demonstrate some of the basic characteristics of the deformable motion algorithm. We will also use two anthropomorphic cardiac phantoms that are designed to provide more complicated deformations more representative of the true motion seen in actual cardiac PET data. These models are now described.

8.1.1. MCAT Phantom

A cardiac PET model based upon an simple geometric structures is one phantom used to test the algorithm. The phantom was obtained using the Mathematical CARDiac Torso (MCAT) software implemented at the University of North Carolina [72,87]. It models the




Figure 8-1. MCAT Cardiac Phantom

Three orthogonal slices through the MCAT cardiac phantom at end diastole and end systole. Organs and other structure within the human torso are modeled using ellipsoidal and cylindrical descriptions. A myocardial defect representing a 20% uptake decrease is seen in the coronal and sagittal views. An isosurface showing the boundaries of the end diastole ventricles is seen on both sets of images.

beating heart inside a fairly realistic torso using a number of ellipsoids and cylindrical building blocks. Effects such as changing myocardial thickness, heart rotation, varying ventricular volume and apical movement during the cardiac cycle are possible with the model. Modeled PET activity for different organs may also be set to a range of different physiologically expected values, including different values for cardiac defects, which typically show up in PET data sets as regions of decreased radiotracer uptake.

Figure 8-1 shows an example of three orthogonal slices through the phantom at end diastole and end systole in a noise-free environment. Voxel intensity for cardiac tissue and other organs in the field of view are set to typical relative PET activities seen in human studies. A myocardial defect is seen in the ventricle wall that represents a 20% decrease in radiotracer uptake. The cardiac activity in the other “healthy” portions of the left ventricle is uniform, and activity in the thinner walls of the right ventricle and atria is set to the same as in the left ventricle. (Note that these thinner walls of the right ventricle and atria are often not even seen in real PET data, due to noise and other effects, as in the example data set




Figure 8-2. Parametric Finite Element Model

Three orthogonal views through a parametric prolate spheroid model of the isolated left ventricle. Sixteen finite elements were used to produce the phantom. The elastic properties of the model were based on MRI data acquired from a canine heart. Included in the model are such factors as tissue incompressibility and cardiac muscle fiber direction.

shown in Figure 3-2). The ratio of voxel intensities between the heart and the blood pool in this data set is 255:6, and between the heart and lung is 255:13.

8.1.2. Prolate Spheroid Finite Element Model

A finite element model of the isolated left ventricle was used to generate a more realistic data set with known motion components. The model is based on a prolate spheroid description of a left ventricle which has been fitted to MRI data acquired from a canine heart [20]. Included in the model is the incompressible nature of cardiac tissue and non-symmetric cardiac muscle fiber orientation. A 16 element model was used to determine the shape of the left ventricle as it was passively inflated. For this paper, the inflated state is used as a reference volume (much like end diastole), and the deflated state is the source volume. Because a parametric description of the two models is available, the “ground truth” motion vectors may be calculated which bring any two points into correspondence. Three orthogonal views of the model are shown in Figure 8-2.




Figure 8-3. Hollow Ellipsoid Model

Three orthogonal views through a hollow ellipsoidal model. The reference volume is seen in the top row. Two versions of a source volume that will be deformed are in the next two rows. The middle row represents a uniform incompressible contraction by a factor of $1/0.8$ and 0.8 along the X and Y axis respectively. The bottom row represents a compressible contraction along the Y axis only by a factor of 0.8 . Isocontours of the reference volume boundaries are seen on all images.

8.2. Material Model Parameter Selection for Isotropic Materials

Though the motion estimation algorithm described in this thesis is designed to model deformations in anisotropic materials, it is instructive to first study the performance of the algorithm for estimating smooth motions generated from an isotropic object. To this end, an isotropic test phantom was designed, which can be seen in Figure 8-3. The phantom is constructed as a hollow ellipsoid with a sinusoidal density pattern on a zero background. Three orthogonal slices through the phantom in its reference configuration are seen in the top row of the figure. An elastic dilation deformation of the model can be easily obtained by changing the scale of the phantom along the X and Y axis. Recall from Chapter 6 that for an incompressible deformation, the product of the extension ratios along the principle

axes must equal 1. That is, $\alpha_1\alpha_2\alpha_3 = 1$. Scaling along the Z-axis was kept constant for these tests. Therefore an incompressible deformation along X and Y could be defined by setting $\alpha_1 = 1/\alpha_2$; and a compressible deformation along Y could be defined by setting

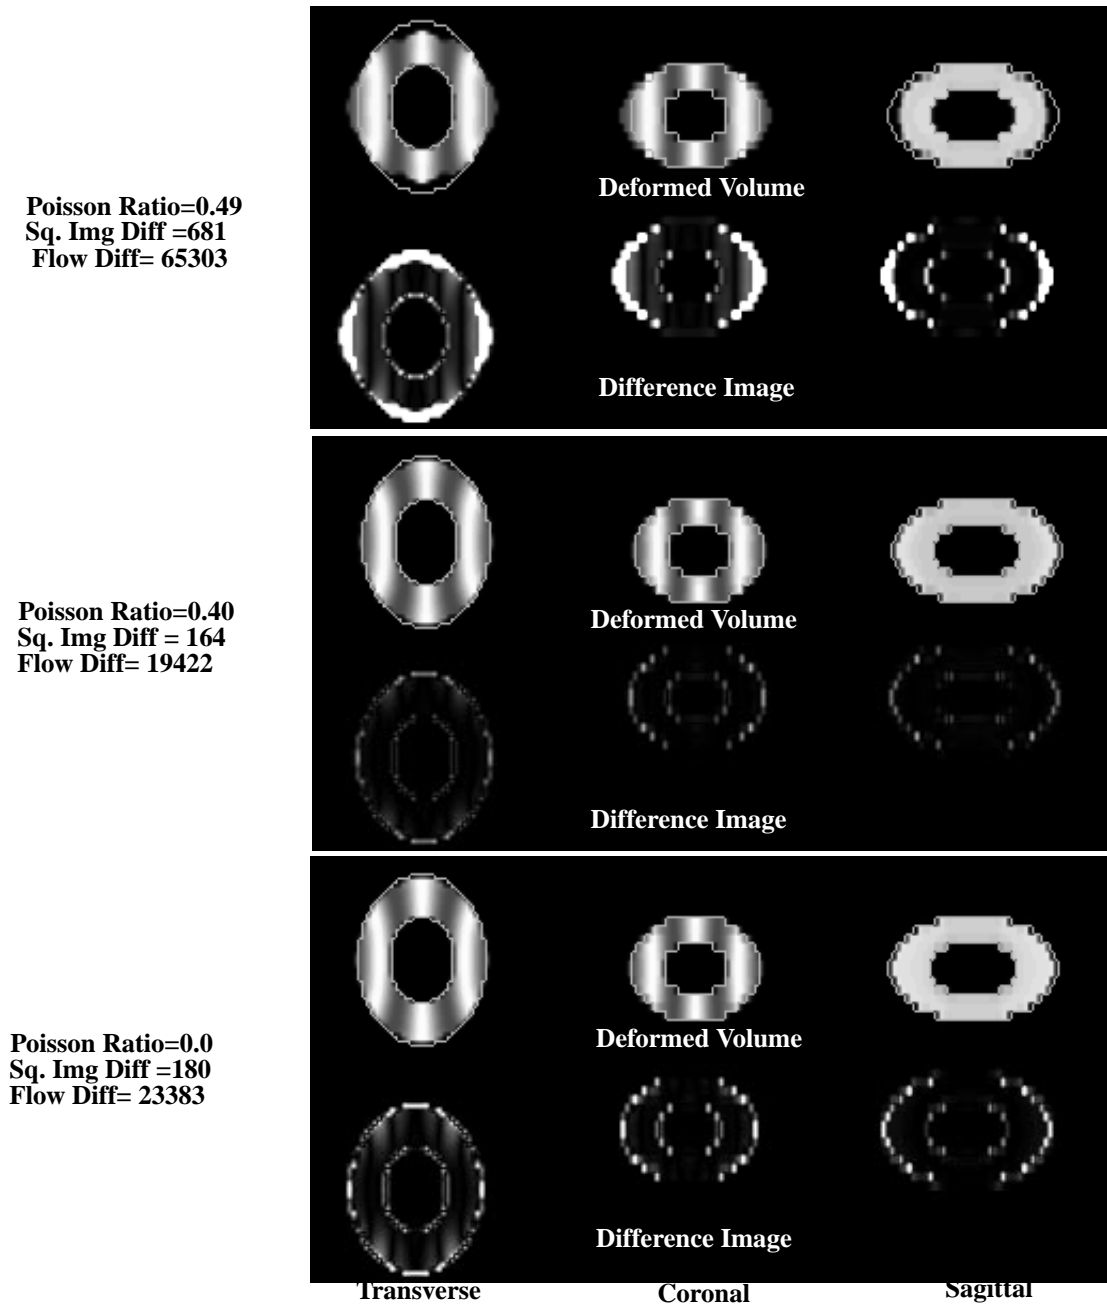


Figure 8-4. Ellipsoid Incompressible Deformation

Deformed volumes and the difference images with respect to the reference volume for modeled Poisson ratios 0.49, 0.40 and 0.0. The best results appear to be obtained using a slightly compressible model even though the phantom was produced using an incompressible assumption.

$\alpha_1 = 1$ and modifying only α_2 . Two different versions of a source volume were produced. The first, seen in the middle row, represents an incompressible expansion along the Y-axis by a factor of 0.8, and a corresponding contraction along the X-axis by a factor of $1.0/0.8$. The second source volume represents a expansion along the Y-axis by the same factor, but with no expansion along the X-axis. The volume size for this phantom is $64 \times 64 \times 16$ with voxel size $1 \times 1 \times 2$ mm.

Results from testing the incompressible deformation are first presented. An isotropic infinitesimal strain energy model with varying values of the Poisson ratio was used by the motion estimation algorithm to compute the motion field. To evaluate the result, the

Poisson Ratio=0.45
Sq. Img Diff =200
Flow Diff=6630




Figure 8-5. Ellipsoid Compressible Deformation

Though the differences are subtle here, using the zero Poisson ratio for a truly compressible deformation produces a better match. Principle differences between the two are around the edges. For high Poisson values, the divergence-free constraint prevented the ellipsoid from fully expanding.

squared difference between the deformed and reference volume was calculated. Additionally, since the true flow field is known in this case, the mean magnitude of the flow vector difference was calculated between the true and estimated flow fields. The resulting deformed volumes for Poisson ratios of 0.49, 0.4 and 0 are shown in Figure 8-4. Comparing the deformed and difference images, it is easy to tell that the result using a very high compressibility factor prevented the algorithm from correctly estimating the result. For lower values of modeled compressibility, the results are more subtle, but the quantitative figures of merit show that the Poisson ratio of 0.4 produced better results. One might expect the model using a Poisson ratio closest to 0.5 to perform best (since we are modeling an incompressible material), however, this is not the result shown in the figure. Instead, the best match appears to be obtained using a slightly compressible model.

This finding is in agreement with results obtained by Gorce, et. al. [36], who investigated the effects of different weightings for the smoothness and divergence-free parameters that were originally proposed by Song and Leahy [82]. They found that ability of the algorithm to accurately estimate motion for an incompressible dilation would decrease when the smoothness penalty was weighted too greatly. This corresponds to using a small or negative Poisson ratio. (Note that a strain energy model with a negative Poisson ratio is possible, though all known materials have Poisson ratios between 0.0 and 0.5.) Gorce obtained their best results when using a largely predominant divergence-free constraint, corresponding to a Poisson ratio close to 0.5.

The effects of model mismatch are further investigated using by estimating the motion field when the actual deformation is a compressible one, as is the depicted in the last row of Figure 8-3. Here, we would expect that the best results would be obtained by choosing a low Poisson ratio which does not penalize regions of non-zero divergence. The expectation is shown to be true in Figure 8-5. It is seen that the motion estimation algorithm works best for Poisson values close to zero. Comparing the images, one can see that in the case of the high Poisson ratio, the divergence-free criteria prevented the source volume from stretching completely along the Y axis to best match the reference.

It is initially a bit perplexing why the algorithm did not make the best motion estimate when using a Poisson ratio closer to 0.5 for the incompressible case. We may shed some light on this result by taking the known motion field and calculating the strain energy as we vary the Poisson ratio. In this case, most of the terms of the motion field gradient are zero so that the infinitesimal strain energy is given by $A = \frac{\lambda}{2}(\alpha_1 + \alpha_2)^2 + \mu(\alpha_1^2 + \alpha_2^2)$. This relation is shown in Figure 8-6b for both the incompressible and compressible cases. This figure points out that even though the motion field was calculated analytically to represent a uniform, incompressible dilation, computation of the strain energy gives a minimum not near 0.5, as expected, but closer to 0.4. In contrast, the dotted line on the same plot shows that for a truly compressible motion field, the minimum strain energy is found for a Poisson ratio of 0.0. Using this line of reasoning, we would likewise expect the deformable motion model to perform best when the modeled Poisson ratio matched the point at which the strain energy for the true flow field minimum exists. That is, for the incompressible deformation in this case we would expect the motion estimation algorithm to perform best with a chosen Poisson ratio near a value of 0.4, and for the compressible deformation, we would expect the best result using a Poisson value of zero.

To see if this is the case, the motion estimation algorithm was run using a wide range of Poisson ratios between 0.0 and 0.5. The results are seen in Figure 8-6c, where the squared difference between the deformed and reference volumes is plotted versus the Poisson ratio used by the estimation algorithm. Again the solid line represents the behavior of the algorithm for the incompressible case. Though the minimum is not exactly at the same place as the analytical strain energy minimum, it is fairly close at a value of about 0.3. Conversely, when the true deformation is compressible, the deformation algorithm performs better as the Poisson ratio approaches zero, as is shown in the dotted line of Figure 8-6c. Note that the analytical curve is fairly flat in (b), so the mismatch does not seem to greatly affect the results until the Poisson ratio grows beyond a value of about 0.3.

So it appears that at least in this simple case, modeling the material as either too incompressible or too compressible produces inferior results. Later results will show that this is true for the more complicated cardiac phantoms as well.

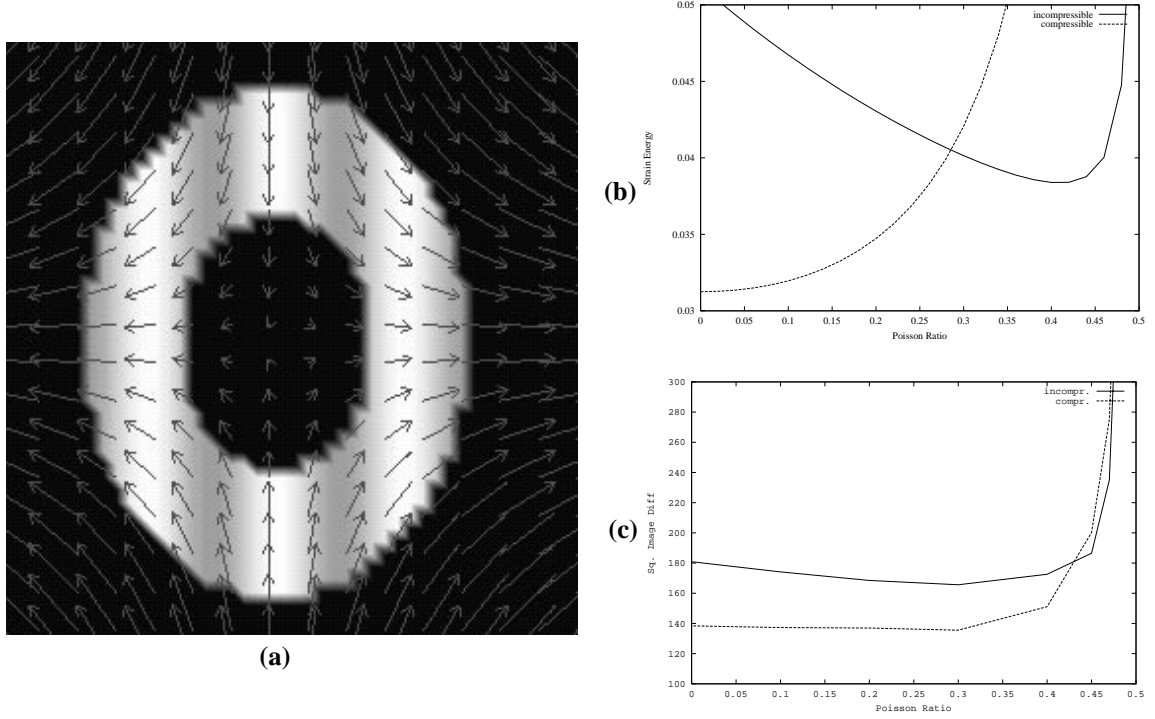


Figure 8-6. Strain Energy for Incompressible Stretch

Incompressible flow field (a), and the analytically calculated strain energy versus Poisson ratio (b - solid line). The minimum strain energy is closer to 0.4 rather than the expected 0.5. This may explain why the motion estimation algorithm performs best using a similar Poisson ratio (c - solid line). In contrast, for the compressible deformation, (dashed line in (b) and (c)), the minimum analytical strain energy exists at zero. Likewise the motion estimate algorithm performs best for a zero Poisson ratio.

8.2.1. Convergence Properties

Even though the displacement vectors were fairly large with respect to the voxel size for the simulations using the ellipsoidal phantom, convergence was possible without using a multi-resolution approach. Recall that the cost function minimization procedure consists of two loops: an outer loop, where a Taylor series linear approximation is made of the deformed volume, and then an inner loop, which uses a conjugate gradient approach to minimize the error criterion with respect to the linearly approximated deformation field. In Appendix A, we showed how the cost function could be expressed in terms of a linear equation in matrix form, $A\mathbf{x} = \mathbf{b}$. A conjugate gradient solution to a set of linear equations like this is guaranteed to converge in at most M iterations, where M is equal the rank

of A [81]. Here, even though A is a $N_1 N_2 N_3$ by $N_1 N_2 N_3$ matrix for a volume with dimensions, N_1, N_2, N_3 , it is found that the algorithm converged relatively quickly, as shown in Figure 8-7. The plot on the left shows the convergence of the inner loop. In this loop, the current motion field has been used to calculate the current deformation volume, and the algorithm calculating the best delta motion field which minimizes the cost function with respect to the Taylor approximation of the deformation field. Then in the outer loop, this delta motion field is added to the current motion field, and a new deformation volume is computed and the process repeats. We see in the right graph that the outer loop converges very quickly; however each time the delta flow field is added, it is seen that the residual error in the inner loop jumps up, then eventually converges to a lower value. Typical processing time to compute 20 outer iterations for the $64 \times 64 \times 16$ voxel ellipsoid data sets was 400 seconds on an SGI R10000 processor running at 225 MHz.

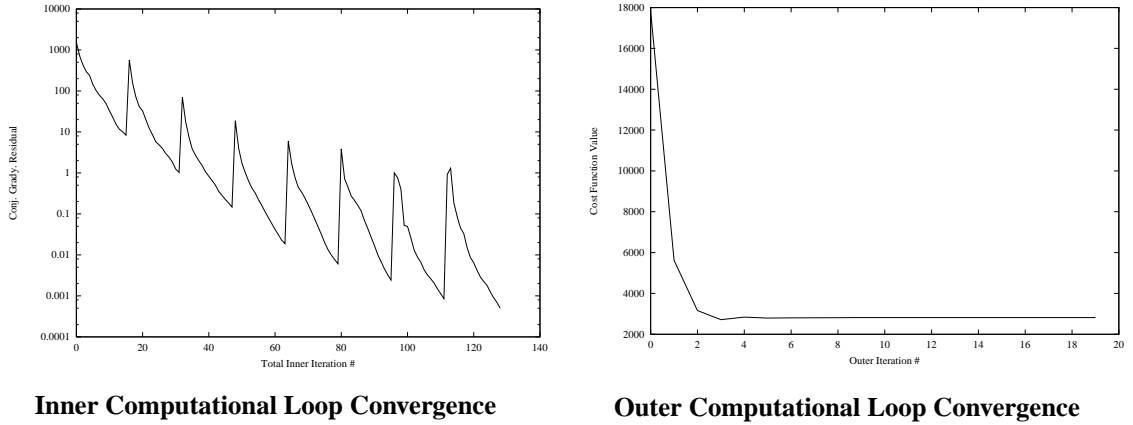


Figure 8-7. Algorithm Convergence Properties

A conjugate gradient minimization implementation of the inner loop usually converged within 5-15 iterations. The outer loop also converged relatively quickly, typically requiring some 5-20 iterations.

8.3. Small Displacement Verses Large Displacement Model Effects

Thus far, the motion estimation results presented on the ellipsoid phantom have all used the infinitesimal strain energy model (small displacement). It turns out that for this case of simple expansion along one axis, and contraction along another, a motion field gradient results with few cross terms. For both the infinitesimal and finite displacement strain

energy models, the only non-zero strain components are the diagonal components, e_{11} and e_{22} . We should therefore not expect a great difference between the results of the motion estimation algorithm using the finite strain energy model or the infinitesimal model, and in fact, this was found to be the case. However, it has been demonstrated that for typical deformations of the heart, the infinitesimal approximation induces large errors in strain estimates. It is yet to be determined as to whether or not these inaccuracies in *strain values* are important for a *strain energy*-based regularization used by a motion estimation algorithm.

To study this question, we will utilize the more complex deformations exhibited by the MCAT and prolate spheroid phantoms. For this experiment, each phantom was still modeled as a single isotropic object, but the regularization constraint was allowed to be either the small displacement strain energy from Equation 6-18, or the finite displacement strain energy defined in Equation 6-19.




Figure 8-8. Strain Energy Model Comparison

The results of the motion estimation algorithm on the MCAT phantom using the infinitesimal and the small displacement strain energy models are shown here. For all chosen values of the Poisson ratio, the finite displacement model appears to perform better.




Figure 8-9. Strain Energy Model Comparison - FEM Phantom

Just as was the case for the MCAT phantom, the finite displacement strain energy model allowed the motion estimation algorithm to obtain results superior to results obtained with the infinitesimal strain energy function. Both the estimated motion field and deformed volumes are less in error using the finite displacement model.

Because the conjugate gradient technique could not be used to minimize the finite displacement version, the nonlinear successive over-relaxation (NLOR) technique was used to minimize the cost function for both strain energy models. It was found that this minimization technique was more prone to local minima, especially when initialized with a motion field of zero and when large deformations were present. This problem occasionally persisted even when using a multi-resolution approach. To solve this issue, the conjugate gradient technique and the infinitesimal strain model was first used to obtain an initial motion estimate, since this technique was found to be much less prone to local minima. The initial estimate was then used as an starting condition for the NLOR technique. Using this hybrid minimization technique produced reliable results more quickly.

Results using the MCAT phantom may be seen in Figure 8-8. For each strain energy model, the motion estimation algorithm was run for a range of Poisson ratios between zero and 0.5. Again, the figure of merit that was used to compare the results was the sum of squared differences between the reference volume and the warped volume. For all the chosen values of the Poisson ratio, the finite displacement strain energy model is superior.




Figure 8-10. Sampling Effects

Deformed MCAT Images (noiseless case). Deformed source volume using forward (a) versus backward sampling (b) are both warped to match the reference quite well, but voxel intensities are closer to being conserved in the forward sampling case. Comparison of the flow magnitude for the forward (c) and backward sampling (d) shows that the backward sampling technique permits concentration of motion along the edges despite an incompressibility constraint in the formulation (see arrow). This results in a less uniformly sampled source image. Note that the overlay of the source edge map on the forward magnitude image(e) and the edge map of the reference image on the backward magnitude (d) show non-zero motion outside the cardiac boundaries, but since the outside voxels are zero-valued, they do not adversely affect the warp.

A similar result is obtained if we use the prolate spheroid finite element model. For this case we compare both the estimated and known flow fields as well as the difference between the deformed and reference volumes. Results are seen in Figure 8-9.

8.4. Sampling Effects

The results of the deformation algorithm showing some of the effects of different sampling techniques are shown in Figure 8-10. The motion estimation algorithm was applied to the MCAT phantom using both a backward sampling and the forward sampling technique and a strain energy regularization allowing a moderate amount of compression. The

image matching term was also weighted fairly heavily so that the algorithm would be attempt to match the deformed volume to the reference even if the motion field because somewhat physically implausible. Both cases show that the heart at systole in the source volume can be warped to result in a image matching the reference image at diastole fairly well. The important distinction here is that the normalized forward mapping technique comes considerably closer to conserving voxel intensities of the source volume. If the grey level of all voxels in the volume are summed, it is seen that the volumes displayed in Figure 8-10a come close to conserving the sum (reference grey sum = 9091, deformed volumes a sum = 8995), whereas for the backward sampled volume Figure 8-10b the difference is larger (voxel grey sum = 8718). Additionally, a close examination of the corresponding motion magnitude image, (Figure 8-10d, f), shows that despite an incompressibility constraint in the matching criteria, the backward sampling technique allows a greater concentration to be sampled from the edges of the ventricle walls, and allows the central portion of the walls to be undersampled. In contrast, the flow magnitude for the forward sampling case (c,e), shows a more uniform appearance indicating that the voxels from the source image are more uniformly distributed into the deformed volume.

The figure also points out in that the highest flow magnitude for the forward sampling occurs in the portion of the volume corresponding to the shape of the heart in the source (e), whereas for the backward sampling, the motion magnitude is effectively showing the motion at the vector “heads” and are thus concentrated at the voxels corresponding to the heart shape in the reference volume. Note that for both the forward and backward sampling cases, regions outside the boundaries of the heart have non-zero motion magnitudes. Though these vectors may not represent true motion of material adjacent to the cardiac walls, the voxel values of this material in PET imagery are close to zero, so they do not adversely affect the deformed image.

8.5. Piece-wise Isotropic Strain Energy Results

It has been demonstrated in the results thus far that by more closely modeling the material properties for an isotropic material, the motion estimation algorithm can better estimate the true motion field. Since we know that cardiac tissue is nearly incompressible, it would

seem that the best results could be obtained using a relatively high Poisson ratio. However, this is complicated by the fact that the total voxel intensity in the region of the blood pool is not conserved. We may actually obtain worse results by modeling the entire volume as an incompressible one. Results using the prolate spheroid model confirm this. Figure 8-11 show the magnitude of the true motion field along with the estimated motion field magnitude for the case where the whole volume was considered fairly incompressible (Poisson ratio=0.45), and where the whole volume was considered very compressible (Poisson ratio=0.1). It is seen that because the area inside the cardiac phantom “cup” is not conserved, the incompressible model attempts to create a large flow magnitude for material exiting the cup, which is in conflict with the true motion field. Conversely, when the incompressibility penalty is relaxed, then the estimated motion is more correct.




Figure 8-11. Motion Flow Error Due to Material Mismatch

Though the cardiac tissue in the prolate spheroid phantom was modeled as incompressible, the blood pool is highly divergent. Therefore using an isotropic incompressible material model in the motion estimation algorithm produced large errors in the motion field estimate. This is especially apparent at the opening of the phantom “cup.”




Figure 8-12. Isotropic vs. Piece-wise Isotropic Model - Noise Free

Assuming the entire volume is a piece of isotropic incompressible media can lead to large errors in the estimated motion when portions of the field do not satisfy this assumption. Using a piece-wise model where the cardiac tissue as modeled as incompressible, and the background as very compressible gives better results.

Therefore it would seem that if we wish to effectively model the material properties of the heart, we need at least two different material assumptions, a fairly incompressible model for cardiac tissue, and a compressible model for the blood pool and possibly other adjacent tissue. This was carried out by using the segmented source volume for the prolate spheroid phantom as a mask, and generating a variable volume for each of the Lamé constants, λ and μ . A comparison of the results for this phantom with the isotropic and piece-wise isotropic models is seen in Figure 8-12. Differences in the image quality are difficult

**Isotropic
Deformed Image
Sq. Error=3111**

**Isotropic
Flow Magnitude
Sq. Error= 22543**

**Piece-Wise Isotropic
Deformed Image
Sq. Error= 2825**

**Piece-Wise Isotropic
Flow Magnitude
Sq. Error = 17149**




Figure 8-13. Isotropic vs. Piece-wise Isotropic - Noisy Case

When noise is added to the prolate spheroid model, both models are able to obtain reasonable estimates of the deformed volume which match the reference shape. Just as in the noise-free case, the piece-wise isotropic results are superior, especially when comparing the squared magnitude difference of the estimated and true motion fields.

to perceive for comparing the two deformed volumes; however, there is a dramatic difference in the accuracy of the estimated motion fields. The quantitative difference measures in the figure represent the squared image difference between the deformed and reference volumes, and the squared error between the magnitude difference of the true and estimated flow field. A similar behavior is seen when we compare the results of the prolate spheroid phantom where additive Gaussian noise has been added to both the source and reference volumes. Seen in Figure 8-13, both the accuracy of the deformed volume and the estimated

**Isotropic
Deformed Image
Sq. Error=2192**

**Piece-Wise Isotropic
Deformed Image
Sq. Error= 843**

**Noisy Case
Isotropic
Deformed Image
Sq. Error= 3415**

**Noisy Case
Piece-wise Isotropic
Deformed Image
Sq. Error= 1896**




Figure 8-14. MCAT - Isotropic vs. Piece-wise Isotropic

Results for both the noisy-free and noisy cases using the MCAT phantom reveal that the deformed error obtained using the piece-wise isotropic model are considerably better than the uniform model. Modeling the divergent blood pool as incompressible appears to be especially harmful when the divergent region has closed boundaries, as is the case here. Dotted lines represent the edges of the reference volume.

motion field are better using the anisotropic model. Here the image difference measure was made with respect to the original (noise-free) reference volume.

The MCAT phantom also shows improvement using the piece-wise isotropic model. Again, for this case, the segmented cardiac tissue in the source volume was used as a mask to define the variable volume of Lamé constants. For the isotropic model, a uniform Poisson ratio of 0.45 was used; for the piece-wise isotropic model, a Poisson ratio of 0.45 was used for the cardiac tissue only, and the background ratio was set to zero. The results are presented in Figure 8-14. Again the quantitative figure of merit used to judge the results is

the squared difference between the deformed and reference volumes. The results show even more dramatically than did the prolate spheroid model that the anisotropic model produced superior results. Perhaps this is due to the fact that the blood pool in this case is completely enclosed by the cardiac tissue. Therefore, trying to use an isotropic incompressible model for this region was even more damaging, as the restraining forces made it difficult for the source volume blood pool to expand to match the reference volume.

A realistic noisy version of the data set was also created by first projecting the model to form noise-free sinograms, and then using these perfect sinograms to obtain a set of statistically sampled sinograms with Poisson statistics, as is standard with simulations of PET data. (note: Poisson *random variables* are unrelated to the Poisson *ratio* discussed thus far). These noisy sinograms were then reconstructed to form a set of noisy images. It is seen in the same figure that when using this noisy data set, the algorithm is still successful in estimating the motion, and that the anisotropic version is superior to the isotropic model.

8.6. Motion Compensation Improvement

The ultimate application for the motion estimation algorithm described in this dissertation is not to obtain the motion field for its own sake, but rather is that it can be used to create a composite data set with less motion blur. Figure 8-15 gives an example of the principle of the motion compensation, and its desired effect. In the top row, we see the noisy version of the MCAT phantom at end diastole. This is the volume we will call the reference. The source volume is the phantom at end systole shown in the second row. The next row represents the result of summing the two volumes without motion compensation. Obviously, because of the dramatic shape changes between diastole and systole, the result of this addition is a degraded data set with considerable motion blur. If instead, we first deform the source volume to match the shape of the reference using the anisotropic motion estimation algorithm, and then add the two together, we obtain the motion compensated data set in the bottom row. Here, we have removed most of the motion blur, and because we have effectively doubled the number of PET counts in the composite volume, the contrast to noise ratio is better than in just the reference volume alone.




Figure 8-15. MCAT - Motion Compensation Example

Noisy versions of the end diastole volume (top row) and the end systole volume (2nd row) for the MCAT phantom. If these are added together directly, considerable motion blur results, and the resulting volume is of less overall quality than either of the individual images. If instead, the systole volume is first warped to match diastole, and then added, the resulting composite has little motion blur and better contrast to noise properties.

This figure points to an application desired in the analysis of real cardiac PET data. Often the whole point of acquiring a cardiac PET study is to visualize which parts of the heart are healthy, and which parts are not. Regions of damaged cardiac tissue are usually characterized by nonuniform uptake patterns, such as the defect seen in the MCAT phantom. It is seen how the uncompensated (direct) sum of the diastole and systole volumes blurs the resolution of the defect, whereas it is more readily apparent in the motion compensated sum.

8.7. Results using Real Data

Now that the phantoms have demonstrated the fundamental behavior of the motion estimation algorithm using a piece-wise isotropic strain energy model, we may present examples of how it works with actual gated cardiac PET data. Figure 8-16 shows three




Figure 8-16. Human PET Data - Warping End Systole

Two 100 msec time frames from a human gated PET sequence representing the heart shape at end diastole (top) and end systole (2nd row). A thresholded version of the systole volume is used as a mask so that the elasticity of this source volume could be modeled anisotropically. The result of warping the source volume to match the reference is seen in the bottom row. The shape of the heart in this warped volume matches the shape at end diastole, yet the voxels have been derived from a realistic elastic warp of the systole data set.

image slices through two cardiac PET volumes which represent the shape of the heart at end diastole and at end systole. These frames were obtained using 100 msec gating intervals. From the rather dramatic shape difference seen between the top two rows, it is obvious that a direct sum of the two frames would introduce considerable motion blur. To reduce this blur, we declare the end diastole data set the reference volume and warp the systole data set so that the shape of the heart matches the reference in both volumes. A mask volume is first defined using a simple thresholding operation on the systole. Elastic properties within this mask are modeled as highly incompressible (Poisson ratio=0.45). For the other regions in the background of the volume, the Poisson ratio is set to zero. The deformed volume obtained using the motion estimation scheme with anisotropic material assumptions is shown in the bottom row of the figure. It can be seen that the shape of the heart matches the end diastole data set quite well, though all voxels values were derived from the source volume.




Figure 8-17. Human PET Data - Summing Comparison

Summing end diastole with the end systole data set results in a volume much like an ungated cardiac PET acquisition. Blurring is proportional to the total motion. By first warping the systole data set to match the reference volume, however, the data can be recombined with minimal motion blur.




Figure 8-18. Human PET Data - Second Subject

Results using data from a different cardiac study are presented. Again, the end systole data set can be warped to match the shape of the end diastole data set. By adding the motion compensated volume with the reference, noise properties of the composite data set are improved over the reference alone.

Results of summing the cardiac gates with and without motion compensation are seen in Figure 8-17. The top row shows a simple sum of the end diastole and end systole volumes. This volume is similar to the images one would obtain in an ungated cardiac sequence. Because voxels at the same locations in the two volumes do not correspond to the same piece of tissue, the motion blur that is introduced by a direct sum is proportional to the amount of motion present. In contrast, the results of the motion-compensated sum is

seen in the bottom row. Here the systole volume is first warped to match the heart shape at end diastole, and then added. The heart walls appear thinner because of less motion blur, and details within the walls are easier to perceive. Because the voxel sum in this composite motion compensated volume represents a sum from corresponding locations in the heart, the contrast to noise properties are improved.

Similar results are presented from another cardiac study in Figure 8-18. Again, considerable heart wall motion is seen between the two 100 msec frames at end diastole and end systole. Just like in the first example, the systole data set was thresholded to obtain a mask used to describe the piece-wise isotropic elastic model. Results of deforming the systole volume via the motion algorithm are shown in the third row. Note that by comparing the fine details in the images between the first three rows, it is seen that the warped volume matches the shape of the reference, however, the algorithm did not introduce features that could not be “stretched” from the source systole data set. The motion corrected composite volume, shown in the bottom row, therefore is a faithful representation of the complete data set.

It is noted that both the source and reference volumes used in the previous example and shown in Figure 8-16 were smoothed considerably using a standard 3D Gaussian spatial filter. Such smoothing is sometimes necessary for the deformable motion algorithm to produce an accurate estimation of the motion field. Since the desired outcome of the motion compensation algorithm is a highly resolved composite volume, it may be preferable to reduce the amount of spatial smoothing, and rely on the improved voxel statistics and motion blur reduction in the composite volume to attain better contrast to noise properties. One approach to achieving this goal for extremely noisy data sets is to use a smoothed version of the source and reference volumes to obtain an accurate motion field estimate, and then use this motion field on the unsmoothed data sets. In this manner, we obtain a composite volume with little spatial blurring due to the Gaussian smoothing filters, yet most of the blur due to motion has been removed.

An example of this is seen in Figure 8-19. Here, the same end diastole and end systole volumes used in the previous example, but with considerably less smoothing are seen in




Figure 8-19. Human PET Data - Summing With Less Smoothing

For extremely noisy cases, a motion field may be computed using source and reference volumes that have been smoothed considerably. This motion field may then be used to recombine the unsmoothed data sets. By this method, contrast to noise properties may be increased in the composite image with minimal spatial blurring not only due to motion artifacts, but also due to the intrinsic blurring properties of image processing filters.

the top two rows of the figure. Again the end systole data set is warped to match the diastole volume, but this time the motion field which is used is the one that was the result of matching the smoothed versions of the data. The result of summing this warped volume with the reference is shown in the bottom row. It is seen that this motion compensated sum appears less noisy than the diastole volume alone, and it has much less motion blur than would be present if the diastole and systole volumes were summed directly. Also, the composite image is more spatially resolved than the case from Figure 8-17 because the raw data volumes were smoothed much less.

Chapter 9

Conclusions

9.1. Discussion

In the deformation of a volume to match a reference data set, there is always a balance between the weight of the image matching constraints and the regularization constraints. The motion estimation problem based on only an image matching constraint is ill-posed. For any pair of source and reference volumes, there are numerous motion fields that can be used to produce a deformed volume identical to the reference volume. It is therefore the function of the regularization constraints to prevent physically unrealizable motion fields. Past techniques have used general smoothness criteria on the motion field to implement a regularization that would prevent unrealistic warps. A problem with this approach is that it is difficult *a priori* to have any idea what appropriate values for the smoothing parameters should be, and how heavily they should be weighted relative to the image matching constraints. In the deformation of real PET data sets, where considerable statistical noise is present, there is always the danger of weighting the image matching terms too greatly so that uncorrelated “hot spots” in the data sets are matched even though they do not originate from the same segment of cardiac tissue. The motivation for this work was to incorporate a more realistic elastic model into the regularization constraint so that this term could be weighted more heavily, and thus would prevent solutions with physically implausible motion fields.

It was shown that use of a linear elastic material model led to weighting terms in the regularization function that had intuitive meanings. For example, the Poisson ratio is a measure of the material compressibility, and for physical materials this quantity is limited to a very restrictive range of values. When the motion field smoothness and divergence-free criteria of other authors are viewed in terms of this model, it is recognized that some parameter weightings are not physically possible in the real world, and therefore are probably not optimal for use as a regularization constraint for the motion estimation problem.

The results presented in this dissertation using phantom data indicate that a better material model leads to better motion estimation. For isotropic materials exhibiting an incompressible stretch, it was shown that an algorithm using a regularization with little penalty for compressive motion fields was less able to accurately estimate the true motion. Similarly, if the true deformation was compressible, and the motion estimation algorithm made the opposite assumption, then results would be degraded as well.

Some past deformable motion techniques have made use of material models, but few have recognized the facts that biological tissue deforms with large-distance deformations, and that in a typical biomedical 3D volume, there is seldom one tissue present with uniform elasticity properties. This is certainly true in imagery of the heart. The left ventricle deforms considerably during the cardiac cycle, indisputably more than the infinitesimal approximation required by some linearly elastic material properties. Also, within the field of view of a typical cardiac study are not only tissue of the myocardium, but also the blood pool within the ventricles, lung air space and bone. These different materials have drastically different elastic properties. Therefore it is no surprise that the finite displacement strain energy function was found to perform better than the strain energy function using the infinitesimal approximation. It is also not surprising that the piece-wise isotropic model performed better than the isotropic material model given the large difference in material properties.

Compensation for contractile motion blur in cardiac PET was the motivation for development of this deformable motion estimation algorithm. With this application in mind, the forward sampling technique was formulated so that a sum of a reference and a deformed

source volume would represent a contribution from every voxel in the original volumes. The validation results carried out on phantoms and the preliminary results obtained on real cardiac PET data indicate that this technique shows promise for use as a reliable motion compensation method. It is acknowledged, though, that the motion of the heart in actual gated PET data is complex, with noise properties not yet completely modeled. Considerably more testing is required to insure that proper correspondence is made between all voxels before image summing takes place, and that the composite image does not produce misleading artifacts.

The choice to use an optical flow-based method rather than a surface model or other high-level model is based partially on the assumption that there is information present in the voxel data which is lost by preprocessing efforts such as edge detection or other high-level parameterizations. Though deformations using higher level descriptions of the object geometry can be very powerful techniques to obtain accurate volume deformations, their success is extremely dependent on the initial feature extraction step. We desire to delay any step requiring an accurate identification of cardiac landmarks. The only preprocessing step required in the described algorithm is a rough segmentation of the cardiac tissue from background tissue. Recognition of the orientation or shape of the heart or identification of specific features is not required. Absolute accuracy of the segmentation should not be that critical either. It appears that the main effect of introducing a nonuniform elasticity description for the myocardial tissue and blood pool was to allow the compressible blood pool to be described differently so that it would not bias the overall estimated motion field. That is, there just needs to be some provision for motion field divergence somewhere within the space of the blood pool. Segmentation errors of just a few voxels from the absolute myocardial-blood pool boundary should therefore not significantly bias the resulting estimated motion field.

Because a 3D optical flow-based method is fairly expensive in terms of computational requirements, we would not like to rule out in the future the use of simpler regularization parameterizations or an image matching criterion based on a sparser set of features. There may be such parameterizations or features that capture essentially the same information as the elasticity model and voxel-based similarity measure used by the algorithm described

here. If so, it would make sense to use these measures, since they could achieve the same results at considerably less computational expense. At this point, though, we feel it is still fruitful to proceed along with the general framework that the optical flow algorithm provides, and to use it to answer further research questions about deformable modeling for feature tracking in biomedical images.

9.2. Future Work

The results obtained thus far appear to indicate that better motion estimations can be obtained by more accurately modeling the material properties of the objects captured in the image data. A natural extension to the work described in this dissertation is therefore investigation of models that characterize the tissue within the field of view in a cardiac PET study even more accurately. For isotropic media, it was found that the finite displacement material model was superior to the infinitesimal model. This model has not yet been tested in the piece-wise isotropic version of this algorithm. It is likely that this improved modeling of cardiac tissue would improve the motion estimation algorithm, however, it is unknown whether the improvement would be worth the extra computational burden. It is doubtful that the same improvement would be seen as was seen by adding the nonuniform elasticity feature to the algorithm, given the extreme difference in material properties between the non-conserved blood pool and the nearly incompressible cardiac tissue.

Along these same lines of development, it may be fruitful to consider some of the other non-linear strain energy models that have been developed using empirical testing on cardiac tissue. An obstacle to this approach is that it would require considerably more preprocessing and pattern recognition. The nonlinear strain energy models for cardiac tissue that was proposed by Nielsen, et. al. [66] is considerably different at directions normal to the surface of the tissue than in directions transverse. Therefore, not only would the heart need to be accurately segmented, but the orientation of the myocardial tissue would also require recognition. Incorporation of additional modeling information, such as the influence of myocardial muscle fiber orientation would require additional recognition as well. Not only would the cardiac tissue have to be identified, and the orientation determined, but the absolute orientation with respect to a pre-defined template would also have to be ascertained.

Accurate characterization of the material is just one facet of possible improvements to this motion estimation algorithm. For motion of tissue due to the contractions of the heart, there is the added advantage that the motion between two time frames is not isolated, but part of a predictable continuous cycle. It is obvious that considerable information is lost by only considering the data in two image volumes when trying to estimate the total motion. The material model used in this dissertation emphasized continuity in the three dimensions of space. The next obvious smoothness constraint to incorporate is smoothness in time. We have investigated one implementation of four-dimensional utilization of information in a previously published paper [48]. In this paper, the information from adjacent time frames was used to calculate a *prediction* motion field, and this motion field was incorporated into the overall cost function for the motion estimation problem. That approach had the one difficulty that in order to provide an initial prediction field, the algorithm needed to be run successfully using only the spatial continuity constraints. Alternate approaches have been proposed that would incorporate four-dimensional continuity constraints more naturally and do not suffer from the initialization problem as does the prediction field method [25, 56,58]. Approaches like these warrant further investigation.

A detailed study of the effects of noise is also needed. Most of the results obtained in this dissertation were obtained using noise-free data. For those datasets, it was advantageous in a motion estimation algorithm to accurately model the elasticity of the material being imaged. This characteristic appears to follow as well for the few noisy cases we tested. On the other hand, a typical method of successfully dealing with noisy data in general has been to apply greater smoothing to the data. Analogously, it may be advantageous to use a material model with slightly “oversmoothed” weighting parameters to obtain better motion estimations for extremely noisy data. The relative weighting of the image matching and regularization constraints also needs to be studied for a wider range of noise characteristics.

One other facet that has been neglected in the modeling of the data thus far is the spatial blurring function associated with the peak spatial resolution of the PET scanner. Given that the cardiac wall thickness varies by almost a factor of two between end diastole and end systole, the apparent thickness of the cardiac walls and voxel intensity in the center of the

walls could be affected considerably in scanners with limited peak spatial resolution. This is a factor that may need to be included in the motion estimation and volume warping calculation.

The results presented in this dissertation have been restricted exclusively to cardiac PET data or phantom data that was emulating characteristics from this particular imaging modality. It may be fruitful to apply the anisotropic approach towards the analysis of MR or CT images of the heart or any other type of data set that could be better modeled using non-uniform material properties. Nearly all the same assumptions that were made regarding the characteristics of PET data sets would apply to voxel data sets from other medical imaging modalities. For example, analysis of motion in data sets of the spine, which combine the rigid motion of the vertebra with the non-rigid motion of adjacent muscle tissue, could benefit from this anisotropic approach. Matching of brain data sets could be enhanced as well, especially in cases where the relative size of the brain ventricles differs radically. Here, just as for the case of modeling the heart ventricles between diastole and systole with an isotropic, incompressible assumption, warping two brains with greatly differing ventricles would be difficult unless a fairly compressible material model was assumed. Better matches might be obtained by modeling the ventricles as compressible material, and the rest of the brain tissue as incompressible material.

Finally, because the ultimate application of this technique is the summing of real cardiac data acquired from human subjects, validation efforts using not only phantom data, but also actual data from clinical PET studies must be carried out. A problem with such validations using real data has always been the difficulty in establishing the ground truth, however, the emergence of other imaging techniques, such as tagged MRI or phase contrast MRI may provide such data sets for accurate evaluation of algorithm performance.

Appendix A

Minimization Via the Calculus of Variations

A.1. Euler-Lagrange Equations

The calculus of variations is a technique useful for determining extreme values of integrals whose integrands consist of unknown functions. For expressions called *functionals* that are functions of functions, the calculus of variations is a procedure for finding a particular function where the functional has a stationary value for small variations about that function. Consider a functional consisting of an integral of the following form:

$$E = \int_a^b F(g, g_x) dx \quad (\text{A-1})$$

where F depends upon an unknown function, $g(x)$ and its derivative, $g_x(x)$. If g is a function that results in the integral at an extrema, then small variations in g will not change the extrema much. More exactly, if we define some test function, $h(x)$, and replace $g(x)$ with $g(x) + \epsilon h(x)$, then the functional will be stationary, that is, it is at an extremum if and only if for all test functions, $h(x)$,

$$\left. \frac{dE}{d\epsilon} \right|_{\epsilon=0} = 0 \quad (\text{A-2})$$

It can be shown [38] that this condition is reached if the following equation, called an *Euler equation*, is satisfied:

$$F_g - \frac{d}{dx}(F_{g_x}) = 0 \quad (\text{A-3})$$

subject to the Dirichlet boundary condition, where we require all test functions, $h(x)$, equal zero on the boundary of volume in which the integral is being evaluated.

Consider now the cost function from Chapter 4-21, expressed here as an integral equation

$$\begin{aligned} E_{tot} = \int_{\Omega} & \left\{ \gamma_I [f_2(\mathbf{r} + \mathbf{m}) - \hat{f}(\mathbf{r} + \tilde{\mathbf{m}}) + \nabla \hat{f}(\mathbf{r} + \tilde{\mathbf{m}}) \bullet \delta \mathbf{m}]^2 + \right. \\ & \left[\frac{\lambda}{2} (\tilde{u}_x - \delta u_x + \tilde{v}_y - \delta v_y + \tilde{w}_z - \delta w_z)^2 + \right. \\ & \mu \left((\tilde{u}_x - \delta u_x)^2 + (\tilde{v}_y - \delta v_y)^2 + (\tilde{w}_z - \delta w_z)^2 \right) + \\ & \frac{\mu}{2} \left((\tilde{u}_y - \delta u_y)^2 + (\tilde{u}_z - \delta u_z)^2 + (\tilde{v}_x - \delta v_x)^2 + (\tilde{v}_z - \delta v_z)^2 + (w_x - \delta w_x)^2 \right. \\ & + (\tilde{w}_y - \delta w_y)^2 + 2(\tilde{u}_y - \delta u_y)(\tilde{v}_x - \delta v_x) + 2(\tilde{u}_z - \delta u_z)(w_x - \delta w_x) \\ & \left. \left. + 2(\tilde{v}_z - \delta v_z)(\tilde{w}_y - \delta w_y) \right) \right] \Bigg\} d\Omega \end{aligned} \quad (\text{A-4})$$

where Ω represents the image volume. Let us call the integrand of this equation, $F(\delta u, \delta v, \delta w, \delta u_x, \delta u_y, \delta u_z, \delta v_x, \delta v_y, \delta v_z, \delta w_x, \delta w_y, \delta w_z)$. Then we may use the multidimensional generalization of the calculus of variation to obtain an extremuum. The cost function may therefore be minimized by solving the following Euler equations:

$$\begin{aligned} F_{\delta u} - \frac{d}{dx}(F_{\delta u_x}) - \frac{d}{dy}(F_{\delta u_y}) - \frac{d}{dz}(F_{\delta u_z}) &= 0 \\ F_{\delta v} - \frac{d}{dx}(F_{\delta v_x}) - \frac{d}{dy}(F_{\delta v_y}) - \frac{d}{dz}(F_{\delta v_z}) &= 0 \\ F_{\delta w} - \frac{d}{dx}(F_{\delta w_x}) - \frac{d}{dy}(F_{\delta w_y}) - \frac{d}{dz}(F_{\delta w_z}) &= 0 \end{aligned} \quad (\text{A-5})$$

Taking for example, the first of these equations, we have

$$\begin{aligned}
F_{\delta u} &= \gamma_I(f_2(\mathbf{r} + \mathbf{m}) - \hat{f}(\mathbf{r} + \tilde{\mathbf{m}}) + \nabla \hat{f}(\mathbf{r} + \tilde{\mathbf{m}}) \bullet \delta \mathbf{m}) \frac{\partial}{\partial x} \hat{f}(\mathbf{r} + \tilde{\mathbf{m}}) \\
\frac{d}{dx}(F_{\delta u_x}) &= \lambda(\delta u_{xx} + \delta v_{xy} + \delta w_{xz} - \tilde{u}_{xx} - \tilde{v}_{xy} - \tilde{w}_{xz}) - 2\mu(\tilde{u}_{xx} - \delta u_{xx}) \\
\frac{d}{dy}(F_{\delta u_y}) &= -\mu(\tilde{u}_{yy} + \tilde{v}_{xy} - \delta u_{yy} - \delta v_{xy}) \\
\frac{d}{dz}(F_{\delta u_z}) &= -\mu(\tilde{u}_{zz} + \tilde{w}_{xz} - \delta u_{zz} - \delta w_{xz})
\end{aligned} \tag{A-6}$$

Similar relations can be found for the last two Euler equations so that the following three partial differential equations already given in Chapter 4 may be obtained:

$$\begin{aligned}
\gamma_I(f_2 - \hat{f} + \nabla \hat{f} \bullet \delta \mathbf{m}) \hat{f}_x &= \lambda(\delta u_{xx} + \delta v_{xy} + \delta w_{xz} - \tilde{u}_{xx} - \tilde{v}_{xy} - \tilde{w}_{xz}) \\
&+ \mu(2\delta u_{xx} + \delta u_{yy} + \delta u_{zz} + \delta v_{xy} + \delta w_{xz} - 2\tilde{u}_{xx} - \tilde{u}_{yy} - \tilde{u}_{zz} - v_{xy} - w_{xz})
\end{aligned} \tag{A-7}$$

$$\begin{aligned}
\gamma_I(f_2 - \hat{f} + \nabla \hat{f} \bullet \delta \mathbf{m}) \hat{f}_y &= \lambda(\delta u_{xy} + \delta v_{yy} + \delta w_{yz} - \tilde{u}_{xy} - \tilde{v}_{yy} - \tilde{w}_{yz}) \\
&+ \mu(\delta v_{xx} + 2\delta v_{yy} + \delta v_{zz} + \delta u_{xy} + \delta w_{yz} - \tilde{v}_{xx} - 2\tilde{v}_{yy} - \tilde{v}_{zz} - u_{xy} - w_{yz})
\end{aligned} \tag{A-8}$$

$$\begin{aligned}
\gamma_I(f_2 - \hat{f} + \nabla \hat{f} \bullet \delta \mathbf{m}) \hat{f}_z &= \lambda(\delta u_{xz} + \delta v_{yz} + \delta w_{zz} - \tilde{u}_{xz} - \tilde{v}_{yz} - \tilde{w}_{zz}) \\
&+ \mu(\delta w_{xx} + \delta w_{yy} + 2\delta w_{zz} + \delta u_{xz} + \delta v_{yz} - \tilde{w}_{xx} - \tilde{w}_{yy} - 2\tilde{w}_{zz} - u_{xz} - v_{yz})
\end{aligned} \tag{A-9}$$

Dirichlet boundary conditions are assumed in the formulation, hence the motion field is assumed to equal zero on the boundaries of the image volume.

A.2. Conjugate Gradient Solution

The partial differential equations expressed in the previous section may be written in matrix form as

$$\mathbf{Ax} = \mathbf{b} \tag{A-10}$$

where \mathbf{A} is a sparse $N_1 N_2 N_3$ by $N_1 N_2 N_3$ matrix for a volume with dimensions, N_1, N_2, N_3 . The matrix is block diagonal with 3x3 blocks as follows:

$$\mathbf{A}_{ii} = \gamma_I \begin{bmatrix} \hat{f}_x \hat{f}_x & \hat{f}_x \hat{f}_y & \hat{f}_x \hat{f}_z \\ \hat{f}_x \hat{f}_y & \hat{f}_y \hat{f}_y & \hat{f}_y \hat{f}_z \\ \hat{f}_x \hat{f}_z & \hat{f}_y \hat{f}_z & \hat{f}_z \hat{f}_z \end{bmatrix} - (\mu + \lambda) \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{xy} & D_{yy} & D_{yz} \\ D_{xz} & D_{yz} & D_{zz} \end{bmatrix} - \mu \begin{bmatrix} D_{xx} & 0 & 0 \\ 0 & D_{yy} & 0 \\ 0 & 0 & D_{zz} \end{bmatrix} \quad (\text{A-11})$$

Likewise, \mathbf{b} and \mathbf{x} are $N_1 N_2 N_3$ column matrices with elements for the i th voxel:

$$\mathbf{b}_i = -\gamma_I (\mathbf{f}_2 - \hat{\mathbf{f}}) \begin{bmatrix} \hat{f}_x \\ \hat{f}_y \\ \hat{f}_z \end{bmatrix} - (\mu + \lambda) \begin{bmatrix} \tilde{u}_{xx} + \tilde{v}_{xy} + \tilde{w}_{xz} \\ \tilde{u}_{xy} + \tilde{v}_{yy} + \tilde{w}_{yz} \\ \tilde{u}_{xz} + \tilde{v}_{yz} + \tilde{w}_{zz} \end{bmatrix} - \mu \begin{bmatrix} \nabla^2 \tilde{u} \\ \nabla^2 \tilde{v} \\ \nabla^2 \tilde{w} \end{bmatrix} \quad \mathbf{x}_i = \begin{bmatrix} \delta u \\ \delta v \\ \delta w \end{bmatrix} \quad (\text{A-12})$$

and where D_{ij} represents a differential operator implemented using a standard finite differencing scheme.

A conjugate gradient scheme is a powerful minimization technique for linear equations in the form of Equation A-10. That is, if we want to minimize

$$\frac{1}{2}(\mathbf{A}\mathbf{x} - \mathbf{b}) = \frac{1}{2}\mathbf{x}^T \mathbf{A}^T \mathbf{A}\mathbf{x} - \mathbf{A}^T \mathbf{b}\mathbf{x} + \mathbf{b}^T \mathbf{b} = \frac{1}{2}\mathbf{x}^T \mathbf{H}\mathbf{x} - \mathbf{e}\mathbf{x} + \mathbf{c} \quad (\text{A-13})$$

then the following conjugate gradient recursion can be used to find the minimum [81]:

	$\mathbf{x} = \text{start}$
Initialization	$\mathbf{r}_0 = \mathbf{e} - \mathbf{H}\mathbf{x}_0$ $\mathbf{p}_0 = \mathbf{r}_0$
<div style="display: flex; align-items: center;"> <div style="margin-right: 10px;">Loop</div> <div style="border-left: 1px solid black; height: 150px; position: relative;"> <div style="position: absolute; top: 0; right: -10px;">→</div> <div style="position: absolute; bottom: 0; left: -10px;">←</div> </div> </div>	$\lambda_n = \frac{(\mathbf{r}_n^T \mathbf{p}_n)}{(\mathbf{p}_n^T \mathbf{H} \mathbf{p}_n)}$ $\mathbf{x}_{n+1} = \mathbf{x}_n + \lambda_n \mathbf{p}_n$ $\mathbf{r}_{n+1} = \mathbf{r}_n + \lambda_n \mathbf{H} \mathbf{p}_n$ $\alpha_n = \frac{(\mathbf{r}_{n+1}^T \mathbf{H} \mathbf{p}_n)}{(\mathbf{p}_n^T \mathbf{H} \mathbf{p}_n)}$ $\mathbf{p}_{n+1} = \mathbf{r}_{n+1} + \alpha_n \mathbf{p}_n$

Furthermore, when A is symmetric, positive-definite, the quadratic form is unnecessary, and we may replace H with A in the recursion. In that case, the norm of the residual, $\|r_n\|$, which now represents the error we are trying to minimize, is guaranteed to decrease.

Bibliography

- [1] A A Amini, Y Chen, R W Curwen, V Mani, and J Sun. "Coupled B-Snake grids and constrained thin-plate splines for analysis of 2-D tissue deformations from tagged MRI." *IEEE Trans Med Imag*, 17(3):344–356, 1998.
- [2] A A Amini and J S Duncan. "Bending and stretching models for lv wall motion analysis from curves and surfaces." *Image and Vision Computing*, 10(6):418–430, 1992.
- [3] Y Amit, U Grenander, and M Piccioni. "Structural image restoration through deformable templates." *J. American Statistical Association*, 86(414):376–387, 1991.
- [4] J L Andersson and L Thurfjell. "Implementation and validation of a fully automatic system for intra- and interindividual registration of PET brain scans." *J Comput Assist Tomogr*, 21(1):136–144, 1997.
- [5] R J Atkin and N Fox. *An Introduction to the Theory of Elasticity*. Longman, London, 1980.
- [6] L Axel and L Dougherty. "MR imaging of motion with spatial modulation of magnetization." *Radiology*, 172:349–350, 1989.
- [7] R Bajcsy and S Kovacic. "Multiresolution elastic matching." *Comput. Vision, Graph., Image Proc.*, 46:1–21, 1989.
- [8] E Bardinet, L Cohen, and N Ayache. "Tracking and motion analysis of the left ventricle with deformable superquadratics." *Medical Image Analysis*, 1(2):129–149, 1996.
- [9] G Barletta, M Baroni, R Del Bene, A Toso, and F Fantini. "Regional and temporal nonuniformity of shape and wall movement in the normal left ventricle." *Cardiology*, 90(3):195–201, 1998.
- [10] F L Bookstein. "Principal warps: Thin-plate splines and the decomposition of deformations." *IEEE Trans Patt Anal Machine Intell*, 11(6):567–585, 1989.

- [11] F L Bookstein. "Shape and the information in medical images: A decade of the morphometric synthesis." *Comput. Vision Image Understanding*, 66(2):97–118, 1997.
- [12] C Broit. *Optimal Registration of Deformed Images*. PhD thesis, University of Pennsylvania, Philadelphia, Pennsylvania, 1981.
- [13] C W Chen, T S Huang, and M Arrott. "Modeling, analysis, and visualization of left ventricle shape and motion by hierarchical decomposition." *IEEE Trans Patt Anal Machine Intell*, 16(4):342–356, 1994.
- [14] G E Christensen, M I Miller, and M W Vannier. "Individualizing neuroanatomical atlases using a massively parallel computer." *Computer*, 29(1):32–38, 1996.
- [15] G E Christensen, R D Rabbitt, and M I Miller. "3d brain mapping using a deformable neuroanatomy." *Phys Med Biol*, 39:609–618, 1994.
- [16] L D Cohen and I Cohen. "Finite-element methods for active contour models and balloons for 2-D and 3-D images." *IEEE Trans Patt Anal Machine Intell*, 15(11):1131–1147, 1993.
- [17] D L Collins, P Neelin, T M Peters, and A C Evans. "Automatic 3D intersubject registration of MR volumetric data in standardized talairach space." *J Comput Assist Tomogr*, 18(2):192–205, 1994.
- [18] R T Constable, K M Rath, A J Sinusas, and J C Gore. "Development and evaluation of tracking algorithms for cardiac wall motion analysis using phase velocity MR imaging." *Magn Reson Med*, 32:33–42, 1994.
- [19] K D Costa, P J Hunter, J M Rogers, J M Guccione, L K Waldman, J M Guccione, and A D McCulloch. "A three-dimensional finite element method for large elastic deformations of ventricular myocardium: I - cylindrical and spherical polar coordinates." *Trans. ASME*, 118:452–463, 1996.
- [20] K D Costa, P J Hunter, J S Wayne, L K Waldman, J M Guccione, and A D McCulloch. "A three-dimensional finite element method for large elastic deformations of ventricular myocardium: II - prolate spheroidal coordinates." *Trans. ASME*, 118:464–472, 1996.
- [21] R Dann, J Hoford, S Kovacic, M Reivich, and R Bajcsy. "Evaluation of elastic matching system for anatomic (CT, MR) and functional PET cerebral images." *J Comput Assist Tomogr*, 13(4):603–611, 1989.
- [22] C Davatzikos. "Spatial normalization of 3D brain images using deformable models." *J Comput Assist Tomogr*, 20(4):656–665, 1996.

- [23] C Davatzikos, J L Prince, and R N Bryan. "Image registration based on boundary mapping." *IEEE Trans Med Imag*, 15(1):112–115, 1996.
- [24] M H Davis, A Khotanzad, D P Flamig, and S E Harms. "A physics-based coordinate transformation for 3-D image matching." *IEEE Trans Med Imag*, 16(3):317–328, 1997.
- [25] J Declerck, J Feldmar, and N Ayache. "Definition of a 4D continuous polar transformation for the tracking and the analysis of LV motion." Technical report, Report 2105, INRIA Sophia-Antipolis France, 1996.
- [26] J Declerck, G Subsol, JP Thirion, and N Ayache. "Automatic retrieval of anatomical structures in 3D medical images." In *Computer Vision, Virtual Reality and Robotics in Medicine, Lecture Notes in Computer Science*, volume 905, pages 153–162. Springer-Verlag, INRIA, Sophia-Antipolis, France, 1995.
- [27] J M Fitzpatrick. "The existence of geometrical density-image transformations corresponding to object motion." *Comput. Vision, Graph., Image Proc.*, 44:155–174, 1988.
- [28] M A Fogel, P M Weinberg, K E Fellows, and E Hoffman. "Single right ventricles compared with systemic right ventricles in a dual-chamber circulation." *Circulation*, 92(2):219–230, 1995.
- [29] M A Fogel, P M Weinberg, K B Gupta, J Rychik, A Hubbard, E Hoffman, and J Haselgrove. "Mechanics of the single left ventricle: A study in ventricular-ventricular interaction II." *Circulation*, 98(4):330–338, 1998.
- [30] P T Fox, J S Perlmutter, and M E Raichle. "A stereotactic method of anatomical localization for positron emission tomography." *J Comput Assist Tomogr*, 9(1):141–153, 1985.
- [31] K J Friston, J Ashburner, C D Frith, J B Poline, J D Heather, and R Frackowiak. "Spatial registration and normalization of images." *Human Brain Mapping*, 2:165–189, 1995.
- [32] K J Friston, C D Frith, P F Liddle, and R Frackowiak. "Plastic transformation of PET images." *J Comput Assist Tomogr*, 15(4):634–639, 1991.
- [33] K P Gallagher, G Osakada, M Matsuzaki, M Miller, W S Kemper, and J Ross. "Nonuniformity of inner and outer systolic wall thickening in conscious dogs." *Am J Physiol.*, 249:H241–248, 1985.
- [34] J Gee, M Reivich, and R Bajcsy. "Elastically deforming a 3D atlas to match anatomical brain images." *J Comput Assist Tomogr*, 17(2):225–236, 1993.

- [35] S Ghosal and P Vanek. "A fast scalable algorithm for discontinuous optical flow estimation." *IEEE Trans Patt Anal Machine Intell*, 18(2):181–194, 1996.
- [36] J Gorce, D Friboulet, and I E Magnin. "Estimation of three-dimensional cardiac velocity fields: Assessment of a differential method and application to three-dimensional CT data." *Med. Image Anal.*, 1(3):245–261, 1997.
- [37] T Greitz, C Bohm, S Holte, and L Eriksson. "A computerized brain atlas: Construction, anatomical content, and some applications." *J Comput Assist Tomogr*, 15(1):26–38, 1991.
- [38] B K P Horn. *Robot Vision*. The MIT Press, Cambridge, Massachusetts, 1986.
- [39] BKP Horn and B G Schunck. "Determining optical flow." *Artificial Intell.*, 17:185–203, 1981.
- [40] R H Huesman, G T Gullberg, W L Greenberg, and T F Budinger. "Users manual: Donner algorithms for reconstruction tomography." Technical Report PUB 214, Lawrence Berkeley Laboratory, 1977.
- [41] R H Huesman, G J Klein, B W Reutter, J H Reed, S Grootoonk, and T F Budinger. "List mode data acquisition for retrospective respiratory-cardiac gated pet." *J Nucl Med*, 39(5 suppl):93P, 1998. (abstract).
- [42] J D Humphrey and F C Yin. "Constitutive relations and finite deformations of passive cardiac tissue II: Stress analysis in the left ventricle." *Circulation Research*, 65(3):805–817, 1989.
- [43] P J Hunter, M P Nash, and G B Sands. "Computational electromechanics of the heart." In A V Panfilov and A V Holden, editors, *Computational Biology of the Heart*, chapter 12, pages 345–404. John Wiley and Sons, Chichester, 1997.
- [44] S C Hunter. *Mechanics of Continuous Media*. Ellis Horwood Limited, Chichester, England, 1976.
- [45] C Kambhamettu and D B Goldgof. "Curvature-based approach to point correspondence recovery in conformal nonrigid motion." *Comput. Vision, Graphics Image Proc.*, 60(1):26–43, 1994.
- [46] M Kass, A Witkin, and D Terzopoulos. "Snakes: Active contour models." *Int. J. Comput. Vision*, 1:321–331, 1988.
- [47] G J Klein. "Forward deformation of PET volumes using material constraints." In B Vemuri, editor, *Proceedings of the Workshop on Biomedical Image Analysis, June 26-27, Santa Barbara, CA*, pages 64–71. IEEE Computer Society, 1998.

- [48] G J Klein and R H Huesman. "A 3D optical flow approach to addition of deformable PET volumes." In *Proceedings of the IEEE Nonrigid and Articulated Motion Workshop, June 15-16, Puerto Rico*, pages 136–143. IEEE Computer Society, 1997.
- [49] G J Klein, B W Reutter, M H Ho, J H Reed, and R H Huesman. "Real-time system for respiratory-cardiac gating in positron tomography." *IEEE Trans Nucl Sci*, 45(4):2139–2143, 1997.
- [50] G J Klein, B W Reutter, M H Ho, J H Reed, and R H Huesman. "Cardiac gating of transmission data is unnecessary for attenuation compensation of double-gated emission scans." In *1998 IEEE Nuclear Science Symposium Conference Record, Toronto, Canada*, 1998.
- [51] J J Koenderink. "The structure of images." *Biol Cybern*, 50:363–370, 1984.
- [52] G W Lenz, E M Haacke, and R D White. "Restospective cardiac gating: A review of technical aspects and future directions." *Magn Reson Imaging*, 7(5):445–455, 1989.
- [53] J Lessick, S Sideman, H Azhari, M Marcus, E Grenadier, and R Beyar. "Regional three-dimensional geometry and function of left ventricles with fibrous aneurysms." *Circulation*, 84(3):1072–1086, 1991.
- [54] A E H Love. *A Treatise on the Mathematical Theory of Elasticity*. Dover Publications, New York, 1927.
- [55] J B Maintz, P A van den Elsen, and M A Viergever. "Evaluation of ridge seeking operators for multimodality medical image matching." *IEEE Trans Patt Anal Machine Intell*, 18(4):353–365, 1996.
- [56] A Matheny and D B Goldgof. "The use of three- and four-dimensional surface harmonics for rigid and nonrigid shape recovery and representation." *IEEE Trans Patt Anal Machine Intell*, 17(10):967–981, 1995.
- [57] C R Meyer, J L Boes, B Kim, P H Bland, K R Zasadny, P V Kison, K Koral, K A Frey, and R L Wahl. "Demonstration of accuracy and clinical versatility of mutual information for automatic multimodality image fusion using affine and thin-plate spline warped geometric deformations." *Med. Image Anal.*, 1(3):195–206, 1996.
- [58] F G Meyer, T Constable, A J Sinusas, and J S Duncan. "Tracking myocardial deformation using phase contrast MR velocity fields: A stochastic approach." *IEEE Trans Med Imag*, 4(453-465), 15.
- [59] G D Meyer, M C Ziskin, W P Santamore, and A A Bove. "Kinematics of the beating heart." *IEEE Trans Bio Engr*, 27(6):319–329, 1980.

- [60] M Miller, G Christensen, Y Amit, and U Grenander. "Mathematical textbook of deformable neuroanatomies." *Proc. Nat'l Academy of Science*, 90(24):11944–11948, 1993.
- [61] S Minoshima, R A Koeppe, K A Frey, M Ishihara, and D E Kuhl. "Stereotactic PET atlas of the human brain: Aid for visual interpretation of functional brain images." *J Nucl Med*, 35(6):949–954, 1994.
- [62] S Minoshima, R A Koeppe, K A Frey, and D E Kuhl. "Anatomic standardization: Linear scaling and nonlinear warping of functional brain images." *J Nucl Med*, 35(9):1528–1537, 1994.
- [63] S C Moore, P F Judy, and D D Garnic. "Prospectively gated cardiac computed tomography." *Med Phys*, 10(6):846–855, 1983.
- [64] M Moshfeghi, S Ranganath, and K Nawyn. "Three-dimensional elastic matching of volumes." *IEEE Trans Med Imag*, 3(2):128–138, 1994.
- [65] H Nagel and W Enkelmann. "An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences." *IEEE Trans Patt Anal Machine Intell*, 8(5):565–593, 1986.
- [66] P M F Nielsen, I J Le Grice, B H Smaill, and P J Hunter. "Mathematical model of geometry and fibrous structure of the heart." *Am. J. Physiol.*, 29:H1365–H1378, 1991.
- [67] W G O'Dell, C C Moore, W C Hunter, E A Zerhouni, and E R McViegh. "Three-dimensional myocardial deformations: Calculation with displacement field fitting to tagged MR images." *Radiology*, 195:829–835, 1995.
- [68] M A Oliver and R Webster. "Kriging: A method of interpolation for geographical information system." *Intl. J. Geographical Information Systems*, 4(3):313–332, 1990.
- [69] N J Pelc, R J Herfkens, A Shimakawa, and D R Enzmann. "Phase contrast cine magnetic resonance imaging." *Magn. Reson. Q.*, 7:229–254, 1991.
- [70] C. A. Pelizzari, G. T. Chen, D. R. Spelbring, R. R. Weichselbaum, and C. T. Chen. "Accurate three-dimensional registration of CT, PET and/or MR images of the brain." *J Comput Assist Tomogr*, 13(1):20–26, 1989.
- [71] A Pentland and B Horowitz. "Recovery of nonrigid motion and structure." *IEEE Trans Patt Anal Machine Intell*, 13(7):730–742, 1991.
- [72] P H Pretorius, W Xia, M A King, B M W Tsui, T S Pan, and B J Villegas. "Evaluation of right and left ventricular volume and ejection fraction using a mathematical cardiac torso phantom." *J Nucl Med*, 38(10):1528–1535, 1997.

- [73] J L Prince and E R McVeigh. "Motion estimation from tagged MR image sequences." *IEEE Trans Med Imag*, 11(2):238–249, 1992.
- [74] P Radeva, A Amini, J Huang, and E Marti. "Deformable b-solids and implicit snakes for localization and tracking of spamm mri-data." In *Proceedings of MMBIA '96*, pages 192–201, Los Alamitos, California, 1996. IEEE Computer Society Press.
- [75] B W Reutter, G J Klein, K M Brennan, and R H Huesman. "Acquisition and automated 3-D segmentation of respiratory/cardiac-gated PET transmission images." In A Del Guerra, editor, *1996 IEEE Nuclear Science Symposium Conference Record, Anaheim, CA November 2-9*, pages 1357–1361, 1996.
- [76] R S Rivlin. "Large elastic deformations of isotropic materials. I." *Proc R Soc Lond A*, 240:459–490, 1948.
- [77] W J Rogers, E P Shapiro, J L Weiss, M B Buchalter, F E Rademakers, M L Weisfeldt, and E A Zerhouni. "Quantification of and correction for left ventricular systolic long-axis shortening by magnetic resonance tissue tagging and slice isolation." *Circulation*, 84(2):721–731, 1991.
- [78] D Rueckert and P Burger. "Geometrically deformable templates for shape-based segmentation and tracking in cardiac mr images." In *Proc. Intl. Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 83–98. Springer-Verlag, 1997.
- [79] P T Sandor and S W Zucker. "Inferring surface trace and differential structure from 3-D images." *IEEE Trans Patt Anal Machine Intell*, 12(9):833–854, 1990.
- [80] R J Seitz, C Bohm, T Greitz, P E Roland, L Eriksson, G Blomqvist, B Rosenqvist, and B Nordell. "Accuracy and precision of the computerized brain atlas programme for localization and quantification in positron emission tomography." *J Cereb Blood Flow Metab*, 10:443–457, 1990.
- [81] G Sewell. *The Numerical Solution of Ordinary and Partial Differential Equations*. Academic Press, Inc., San Diego, 1988.
- [82] S M Song and R M Leahy. "Computation of 3-D velocity fields from 3-D cine CT images of a human heart." *IEEE Trans Med Imag*, 10(1):295–306, 1991.
- [83] S M Song, R M Leahy, D P Boyd, B H Brundage, and S Napel. "Determining cardiac velocity fields and intraventricular pressure distribution from a sequence of Ultrafast CT cardiac images." *IEEE Trans Med Imag*, 13(2):386–397, 1994.
- [84] L H Staib and J S Duncan. "Boundary finding with parametrically deformed models." *IEEE Trans Patt Anal Machine Intell*, 14(11):1061–1075, 1992.

- [85] K S Sunnerhagen, B Bhargava, and R Shabetai. "Regional left ventricular wall motion abnormalities in idiopathic dilated cardiomyopathy." *Am J Cardiol.*, 65(5):364–370, 1990.
- [86] J Talairach and P Tournoux. *Co-planar Stereotactic Atlas of the Human Brain: 3-Dimensional Proportional System: An Approach to Cerebral Imaging*. Georg Thieme Verlag, Stuttgart, 1988.
- [87] J A Terry, B M W Tsui, J R Perry, and G T Gullberg. "A three-dimensional mathematical phantom of the human torso for use in SPECT imaging research studies." *J Nucl Med*, 31(5):868, 1990.
- [88] D Terzopoulos and D Metaxas. "Dynamic 3D models with local and global deformations: Deformable superquadrics." *IEEE Trans Patt Anal Machine Intell*, 13(7):703–714, 1991.
- [89] D Terzopoulos, A Witken, and M Kass. "Constraints on deformable models: Recovering 3D shape and nonrigid motion." *Artificial Intell.*, 36:91–123, 1988.
- [90] J P Thirion. "Fast non-rigid matching of 3d medical images." In *Medical Robotics and Computer Aided Surgery*, pages 47–54, Baltimore, 1995.
- [91] J P Thirion and A Gourdon. "Computing the differential characteristics of iso-intensity surfaces." *Comput. Vision Image Understanding*, 61(2):190–202, 1995.
- [92] P Thompson and A W Toga. "A surface-based technique for warping three-dimensional images of the brain." *IEEE Trans Med Imag*, 15(4):402–417, 1996.
- [93] L Thurfjell, C Bohm, T Greitz, and L Eriksson. "Transformations and algorithms in a computerized brain atlas." *IEEE Trans Nucl Sci*, 40(4):1187–1191, 1993.
- [94] D L Vine, H T Dodge, M Frimer, D K Stewart, and J Caldwell. "Quantitative measurement of left ventricular volumes in man from radioopaque epicardial markers." *Circulation*, 54:391–398, 1997.
- [95] L K Waldman, Y C Fung, and J W Covell. "Transmural myocardial deformation in the canine left ventricle." *Circulation Research*, 57(1):152–163, 1985.
- [96] K Washizu. *Variational Methods in Elasticity and Plasticity*. Pergamon Press, Oxford, 1982.
- [97] K Weinhard, M Dahlbom, L Eriksson, C Michel, T Bruckbauer, U Pietrzyk, and WD Heiss. "The ECAT EXACT HR: Performance of a new high resolution positron scanner." *J Comput Assist Tomogr*, 18(1):110–118, 1994.

- [98] A Witkin. "Scale space filtering." In A Bundy, editor, *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 1019–1022, 1983.
- [99] R P Woods, S T Grafton, C J Holmes, S R Cherry, and J C Mazziotta. "Automated image registration: I. general methods and intrasubject, intramodality validation." *J Comput Assist Tomogr*, 22(1):139–152, 1998.
- [100] R P Woods, S T Grafton, J D Watson, N L Sicotte, and J C Mazziotta. "Automated image registration: II. intersubject validation of linear and nonlinear models." *J Comput Assist Tomogr*, 22(1):153–165, 1998.
- [101] A A Young and L Axel. "Three-dimensional motion and deformation of the heart wall: Estimation with spatial modulation of magnetization - a model-based approach." *Radiology*, 185:241–247, 1992.
- [102] E Zerhouni, D M Parish, W F Rogers, A Yang, and E P Shapiro. "Human heart: Tagging with MR imaging - a method for noninvasive assessment of myocardial motion." *Radiology*, 169:59–63, 1988.
- [103] Z Zhou, C E Synolakis, R M Leahy, and S M Song. "Calculation of 3D internal displacement fields from 3D X-ray computer tomographic images." *Proc R Soc Lond A*, 449(1937):537–554, 1995.